NON-CONFIDENTIAL
2015-1256

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

Wi-LAN USA, INC., and Wi-LAN, INC.,
Plaintiffs-Appellants,

v.

APPLE INC.,
Defendant-Appellee,

Appeal from the United States Court for the Southern District of California in Case No. 3:13-cv-0798, Honorable Judge Dana Sabraw

NON-CONFIDENTIAL OPENING BRIEF FOR PLAINTIFFS-APPELLANTS, Wi-LAN USA, INC., AND Wi-LAN, INC.

Robert A. Cote
Kevin Schubert
McKOOL SMITH P.C.
One Bryant Park, 47th Floor
New York, New York 10036
Tel: (212) 402-9400
Fax: (21) 402-9444
rcote@mckoolsmith.com
kschubert@mckoolsmith.com

Dirk D. Thomas
McKOOL SMITH P.C.
1999 K Street, Suite 600
Washington, DC 20006
Tel: (202) 370-8302
Fax: (202) 370-8344
DThomas@McKoolSmith.com

Attorneys for Plaintiffs-Appellants,
Wi-LAN USA, INC., and Wi-LAN, INC.

March 13, 2015

Claim 1 of the '040 Patent

1. A node for a communications system that packs and fragments variable-length service data units (SDU) for mapping into variable length protocol data units (PDU), each SDU being **associated with a specified connection**, the node comprising:

a communications processor configured to pack and fragment SDUs associated with **a specified connection** into a PDU, including

allocating bandwidth for **the specified connection**, based on the priority of the connection,

establish a length for the PDU based on the bandwidth allocated to **the specified connection** in a current frame,

pack a first SDU into a payload area of the PDU,

determine whether a second SDU is larger than a remaining payload area of the PDU,

if the second SDU is not larger than the remaining payload area of the PDU, map the second SDU to the remaining payload area of the PDU, and

if the second SDU is larger than the remaining payload area of the PDU, fragment the second SDU into at least two fragments and map the first fragment to the remaining payload area of the PDU, and

including packing sub-headers in the PDU to allow determination of the length of the SDUs and the lengths of the fragments that are mapped to the PDU.

JA55 (19:29-53)

Claim 1 of the ’640 Patent

1. A method for requesting bandwidth on demand in a wireless communication system, wherein the wireless communication system includes a wireless subscriber radio unit, the method comprising:

registering the wireless communication radio unit with a base station in the wireless communication system and establishing communication between the wireless subscriber radio unit and the base station;

transmitting from the wireless subscriber radio unit which is registered with the base station, an explicit message to the base station requesting to be provided an allocation of uplink (UL) bandwidth in which to transmit a bandwidth request;

receiving at the wireless subscriber radio unit the allocation of UL bandwidth in which to transmit a bandwidth request;

transmitting the bandwidth request within the allocation of UL bandwidth, the bandwidth request specifying a requested UL bandwidth allocation; and

receiving an UL bandwidth grant for the wireless subscriber radio unit in response to the bandwidth request;

wherein the wireless subscriber radio unit maintains a plurality of queues, each queue for data pertaining to one or more **UL connections** with similar QoS and wherein the wireless subscriber radio unit allocates the UL bandwidth to the one or more **UL connections** based on QoS priority.

JA84 (23:7-33)

CERTIFICATE OF INTEREST

Counsel for Plaintiffs-Appellants Wi-LAN USA, Inc. and Wi-LAN Inc. certifies the following:

1. The full name of every party represented by me is: Wi-LAN USA, Inc. and Wi-LAN Inc.
2. The name of the real party in interest represented by me is: N/A.
3. All parent corporations and any publicly held companies that own 10 percent or more of the stock of the party represented by me is: N/A.
4. The names of all law firms and the partners or associates that appeared for the party represented by me in the trial court or are expected to appear in this Court are:

McKOOL SMITH P.C.; Robert Cote, Dirk Thomas, Brett Cooper, Jonathan Yim, Seth Hasenour, Kevin Schubert, Lawrence M. Hadley

PATTERSON LAW GROUP; Allison H. Goddard

VINSON & ELKINS LLP; Andrea Houston, Charles Ebertin, Efren Garcia, Seth Adam Lindner, Stephen Hash, David Weaver, Steven Borgman

Dated: March 13, 2015.

Respectfully submitted,

/s/ Dirk Thomas
Dirk D. Thomas
McKOOL SMITH P.C.
1999 K Street, Suite 600
Washington, DC 20006
Attorney for Plaintiffs-Appellants, Wi-LAN USA, Inc., and Wi-LAN Inc.

TABLE OF CONTENTS

Page

TABLE OF AUTHORITIES

Page(s)

CASES

STATUTES

STATEMENT CONCERNING CONFIDENTIAL MATERIAL

Pursuant to FED. CIR. R. 28(d)(1)(B), Plaintiff-Appellants Wi-LAN USA, Inc. and Wi-LAN Inc ("Wi-LAN") state as follow: confidential material id redacted in this brief at pages 3, 10, 12, 13, 15, 24, 25, 28, 33, 34, 35, 36, 37 and at JA6 – JA16 of the Addendum consisting of trade secret information regarding microcips and descriptions of information regarding trade secrets documents filed under seal and protected under the protective order entered by the District Court of the Southern District of California.

STATEMENT OF RELATED CASES

Pursuant to FED. CIR. R. 47.5, Plaintiff-Appellants Wi-LAN USA, Inc. and Wi-LAN Inc ("Wi-LAN") state that there have been no previous appeals in this case. Counsel for Wi-LAN are unaware of other pending cases in this Court that will be directly affected by, or that will directly affect, this Court's decision on the pending appeal, No. 2015-1256.

I. STATEMENT OF JURISDICTION

The district court for the Southern District of California ("district court") had jurisdiction over this action under 28 U.S.C. §§ 1331 and 1338(a) and issued a final judgment of non-infringement on January 8, 2015, as to all asserted claims of U.S. Patent Nos. 8,315,640 ("the '640 patent") (JA57-84) and 8,311,040 ("the '040 patent") (JA29-56). JA2. The Court of Appeals for the Federal Circuit has jurisdiction over this appeal under 28 U.S.C. § 1295(a). Wi-LAN timely filed its notice of appeal on January 12, 2015, under FED. R. APP. P. 4(a) and 28 U.S.C. § 2107(a). JA85-86.

II. STATEMENT OF THE ISSUES

This appeal is narrowly focused on two claim construction issues:

Issue 1: Does the term "UL [uplink] connections" in the asserted claims of the '640 patent refer to connections located between the subscriber unit and a base station (as the patentee expressly defined the "uplink"), or connections between the subscriber unit and an end user (as construed by the district court)?

Issue 2: Does the term "associated with a specified connection" in the asserted claims of the '040 patent require packing and fragmenting SDUs from *one* specified connection into a PDU for transport to a base station (as recited by the plain language), or from *multiple* specified connections (as construed by the district court)?

III. STATEMENT OF THE CASE SETTING OUT THE FACTS RELEVANT TO THE ISSUES

A. Preliminary Statement

Shortly before trial, the district court granted summary judgment of non-infringement based on two claim constructions that contradict the intrinsic evidence. This narrow appeal asks this Court to revisit these two claim constructions and reverse the summary judgment of non-infringement based on the constructions.

First, the district court misconstrued the term "UL connection" in the '640 patent as a connection between a subscriber radio unit and an end user (e.g. the application processor on the mobile phone that generates data). The "UL connection" is, in fact, the wireless connection from the subscriber radio unit to the cellular network base station, consistent with the patentee's definition and the intrinsic evidence. The patentee explicitly defined "uplink" (abbreviated "UL")[1] to be from the subscriber unit to the base station:

> Transmissions **from the subscriber unit to the base station** are commonly referred to as 'uplink' transmissions.

[1] The '640 patent equates "uplink" and "UL." JA84 (23:18) ("uplink (UL)"). As the district court stated, "[t]here is no dispute 'UL' means 'uplink.'" JA27.

JA73 (1:51-52).[2] Instead of applying the patentee's definition of "uplink," the district court chose a different definition and supported that choice with no evidence. In addition to ignoring the patentee's chosen definition, the district court's construction contradicts other intrinsic evidence, including the claim language, the specification, and the prosecution history.

Separately, the district court misconstrued the claim term "associated with a specified connection" in the '040 patent to require that SDUs (e.g. end user application data) being packed and fragmented into a PDU at the subscriber unit (for transport to a base station) are from multiple specified connections, despite the claim language to the contrary, [redacted] The district court's construction also excludes a preferred embodiment: packing data from the "same connection" to conserve bandwidth. JA2120. Wi-LAN asks this Court to construe "associated with a specified connection" according to its plain meaning and require only one specified connection.

B. Overview of the Patents and Problems Solved

This appeal involves two patents, the '640 patent and the '040 patent. Although the '640 patent and the '040 patent share one overlapping inventor, Ken

[2] Throughout this brief, bolding and underlining should be presumed to have added emphasis to the original, unless otherwise noted.

Stanwood, the patents are not related and have completely different specifications. Thus, the two claim construction issues in this appeal must be separately analyzed.

Each patent generally optimizes the use of the wireless "uplink" for transporting data from a subscriber radio unit (or node) to a base station, but through different means. An "uplink" connection is a communication link over which data is wirelessly transported to a base station. JA73 (1:51-52). Numerous cellular devices, such as the accused Apple iPhones, compete for the uplink bandwidth that a particular base station can grant for use in transporting wireless data to the base station on an "uplink" ("UL") connection. JA754. That cellular bandwidth comprises a scarce commodity needed to transmit pending data at the subscriber radio unit to the cellular base station. JA2480 (¶8).

The explosion in the amount of data that mobile devices are transmitting and receiving over cellular networks makes the efficient use of that bandwidth critical because the licensed wireless bandwidth available to transmit that data remains fixed. Figure 1 from the '640 patent (below) shows two cell sites, each with a cellular base station (106) and a plurality of cellular subscriber radio units (110) in buildings communicating with the base station.[3] JA75 (6:3-7).

[3] One type of cellular device is a Customer Premises Equipment (or "CPE"). The claims are directed to a "wireless subscriber [communication] radio unit" ('640 patent) or "node" ('040 patent), which the district court construed consistently. JA18 (fn. 2). For purposes of this appeal, "node" and "wireless subscriber

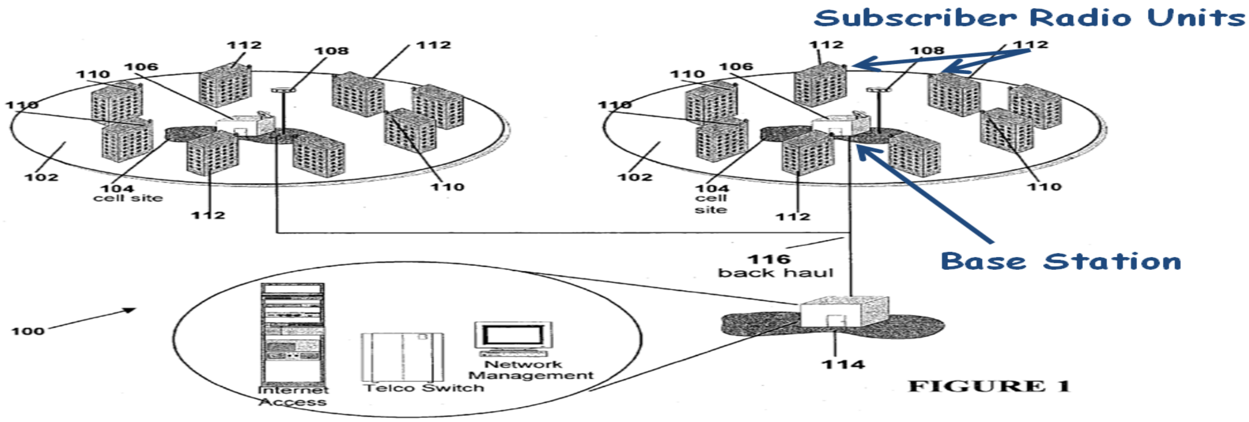

JA60 (Figure 1 (annotated)). The ’640 patent describes the problem of numerous subscriber radio units attempting to coordinate their transmissions to a cellular base station over limited bandwidth:

> [I]n the system shown in Fig. 1, as many as one hundred CPEs [subscriber radio units] may be allowed to be simultaneously active, ***coordinating their transmissions on the uplink***. . . . The bandwidth allocation method and apparatus should accommodate an arbitrarily large number of CPEs generating frequent and varying bandwidth allocation requests on the uplink of a wireless communication system. Such a bandwidth allocation method should be ***efficient in terms of the amount of bandwidth*** consumed by the bandwidth request control messages exchanged between the plurality of base stations and the plurality of CPEs.

JA74 (3:7-44).

As explained in more detail below, each patent-in-suit optimizes the wireless uplink connection between the subscriber radio unit and the base station through

---

[communication] radio unit” are sometimes referred to in short-hand as “subscriber radio unit” or “subscriber unit.” The district court determined that a CPE is a type of subscriber radio unit that is fixed, but that a subscriber radio unit may also be mobile. JA26-27; JA18.

different means: the ’640 patent by efficiently requesting and allocating bandwidth to “UL connections” based on priority so that high priority data (whether end user application data or control data) is sent first, and the ’040 patent by packing and fragmenting such data into a PDU at the subscriber radio unit to most efficiently use the allocated bandwidth so that as much data as possible can be transported to the base station.

1. The ’640 Patent

The ’640 patent provides a two-step bandwidth request and allocation process that minimizes the uplink bandwidth overhead consumed in the process of requesting an uplink bandwidth grant from the base station. JA73 (2:55-66); JA319-320. The ’640 patent also teaches allocating the uplink bandwidth grant to one or more “UL connections” that are established at (i.e. originate at) the subscriber radio unit and transport data to the base station. JA73 (1:49-65); JA74 (4:49-54). Data for higher priority connections (e.g. voice, video, or control messages) is allocated uplink bandwidth before data for lower priority connections (e.g. email or other internet protocol (IP) data), such that the higher priority data is generally transmitted to the base station before the lower priority data. JA76 (7:8-12); JA77 (10:53-56); JA82 (19:29-44).

Figure 10 below shows two uplink (“UL”) connections between the subscriber radio unit (“CPE”) and the base station (“BS”): “connection n” and

"connection k." Each connection may carry, for example, voice, video, control messages or other types of data between the subscriber radio unit and the base station. JA73 (1:62-65); JA79 (14:46-57).

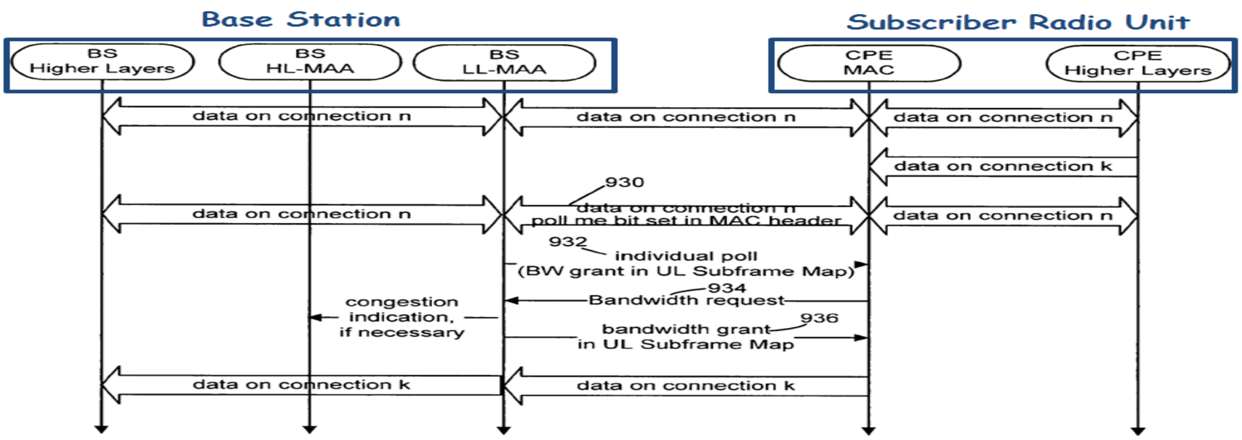

FIGURE 10

JA69 (Fig. 10) (annotated). As illustrated above, the subscriber radio unit ("CPE") allocates bandwidth first to uplink "connection n" for use by higher priority data. While communicating data to the base station ("BS") for "connection n," the subscriber radio unit has lower priority data needing transmission on "connection k." Using a very small amount of the bandwidth allocated to "connection n," the subscriber radio unit sends the base station a "poll me bit" (930, "step 1") to indicate that additional bandwidth is required. JA81 (17:57-61).

The base station grants enough bandwidth (932) to the subscriber unit for it to make a request for bandwidth (934, "step two"). JA81 (17:61-66). Finally, additional bandwidth is granted (936) to the subscriber radio unit which may be allocated to "connection k." JA81 (17:66-18:4).

In sum, the ’640 patent teaches (i) requesting uplink bandwidth in an efficient, two-step manner that minimizes the uplink bandwidth overhead consumed in the request process, and (ii) allocating the bandwidth to one or more “UL connections” between the subscriber radio unit and the base station, such that the highest priority data is sent first. JA83 (22:24-50); JA298-299.

2. The ’040 Patent

Like the ’640 patent, the ’040 patent makes efficient use of the uplink wireless bandwidth that is licensed by the cellular telephone companies and shared by a large number of nodes (subscriber radio units) transmitting data to a single cellular base station. JA46 (1:53-57); JA2120-2121 (describing a savings in uplink bandwidth); JA1852 (“This subject matter allows, *inter alia*, the method to better use the bandwidth shared by a plurality of subscriber stations.”). The ’040 patent, packs and fragments multiple SDUs from the <u>same</u> connection (the “specified connection”) to conserve uplink bandwidth:

> Packing ***multiple SDUs from the <u>same connection</u>*** behind a single MAC header ***saves bandwidth*** over attaching a separate MAC header to each . . . . For instance, the uplink MAC header is more than 16% of the size of an ATM payload.

JA2120-2121.[4] In other words, the ’040 patent efficiently fragments and populates SDUs from the same connection into what is called a PDU, the vehicle which transports the SDUs to the base station, in order to transport as much data as possible within the allocated bandwidth. JA47 (19:41-50, 2:3-8). In one embodiment, the subscriber radio unit may format “data arriving in SDUs of various formats into” a PDU. JA47 (2:3-8). For example, SDUs may arrive at the subscriber radio unit from a user application in a fixed length or in a variable length: “For each system these packets [SDUs] may be of a standard length or may vary in length as the needs of the users dictate.” JA46 (1:29-32).

In sum, regardless of whether the end user application provides SDUs of fixed length or variable length, the subscriber radio unit (or “node”) efficiently packs and fragments as many SDUs as possible from the specified connection into a PDU to conserve uplink bandwidth. JA47 (3:43-55); JA2120-2121.

C. Ensemble and Its Contribution to 4G Technology

The technology in the patents-in-suit was developed at Ensemble, a wireless company in the San Diego area that was founded in 1997 and grew to over 200 employees. JA171. Ensemble was one of the “Hottest Start-ups of 2000,” according to Telecommunications Magazine. JA2579. Wi-LAN and Ensemble had a partnership for developing 4G products in the early 2000s. Beginning in

[4] The quoted passage appears in the provisional application, which is *incorporated by reference in its entirety into the ’040 patent specification*. JA46 (1:7-11).

2004, Ensemble began selling its assets, including the applications leading to the patents-in-suit, as part of winding down Ensemble's operations. JA171.

Ensemble was a major contributor to the 4G technologies in use today. JA171-172. Lead inventor Mr. Stanwood was the Vice-Chair of the committee that standardized the first 4G standard (referred to as "WiMAX" or "IEEE 802.16"[5]) and a principal drafter and technology contributor of that standard. *Id.* The technology and benefits of the patents-in-suit were explained to that committee before being adopted into the 4G standard. *Id.*

Notably, the 4G WiMAX standard is connection-oriented and uses logical connections (referred to as "UL connections" in the '640 patent) to exchange data between a subscriber radio unit and a base station, [redacted]

[redacted]

[redacted] The accused Apple products practice 4G LTE, the most recent adaptation of the 4G cellular standard developed by Mr. Stanwood and others, which incorporates much of the foundational technology from the predecessor 4G WiMAX standard, including the 4G technologies in the patents-in-suit. JA172-181. Mr. Stanwood continues his work as CEO of Wi-LAN Labs (formerly

[5] A draft of this standard was included with the provisional application of the '040 patent. JA2126.

Cygnus Broadband), an arm of Wi-LAN based in San Diego and focused on developing advanced 4G technologies and products. JA167.

D. The Claims At Issue and Accused Products

Wi-LAN asserted independent claims 1 and 6 (and dependent claims 2 and 7) of the ’640 patent in the district court. JA12. Wi-LAN appeals only the summary judgment ruling for independent claim 1 and dependent claim 2 (the "’640 claims"). The term "UL connections" is used in the ’640 claims at issue in this appeal. Wi-LAN asserted independent claims 1, 14, and 16 and dependent claims 2, 4, 5, and 15 of the ’040 patent (the "’040 claims") in the district court and appeals all of these ’040 claims. JA9. The term "associated with a [the] specified connection" is used in all of the ’040 claims.[6]

Wi-LAN has accused Apple’s iPhone and iPad products that operate in accordance with the 4G LTE wireless cellular standard of infringement.

E. The District Court’s Claim Construction Orders

1. *Markman* Order

The district court issued a *Markman* Order on December 23, 2013. For the ’640 patent, the district court provided the following construction for "UL

[6] Independent claim 16 of the ’040 patent uses the term "associated with <u>the</u> specified connection." JA56 (21:27-28). Independent claims 1 and 14 use the term "associated with <u>a</u> specified connection." JA55 (19:31, 20:53).

connections": "an uplink connection between the wireless subscriber radio unit and its users." JA27.

Wi-LAN had proposed that "UL connections" means "uplink services" between the subscriber unit and the base station. JA73 (1:59-65). Wi-LAN argued that the patents-in-suit optimize the uplink, as Wi-LAN has consistently argued throughout the case. JA91-92. The district court did not adopt Wi-LAN's construction, and, instead, adopted Apple's construction nearly verbatim, citing no evidence. JA27.

For the '040 patent, Apple did not suggest that the term "associated with a specified connection" needed construction or argue that it required multiple connections during the *Markman* process. Apple first raised this argument in its summary judgment papers. *Id.*

F. Apple Raises New Claim Construction Issues on Summary Judgment Which the District Court Grants

Apple filed a Motion for Summary Judgment of Non-Infringement on August 22, 2014, arguing that multiple connections are required by the claims of both patents-in-suit [redacted] Specifically, Apple argued more than one "specified connection" and "UL connection" are required by the asserted claims. Apple never raised this claim construction issue during the *Markman* process and first argued these terms require

multiple connections in its summary judgment papers, as the district court acknowledged:

[redacted]

[redacted] The district court construed the asserted claims of both patents-in-suit for the first time in its summary judgment order to require multiple connections and, based on that construction, entered summary judgment of non-infringement. JA11-14.

G. Faced With the District Court's New Constructions, Wi-LAN Seeks Reconsideration

For the '040 patent, Wi-LAN sought reconsideration of the district court's construction that "associated with a specified connection" requires more than one connection, as the district court decided for the first time in its summary judgment order. JA1094-1100. Wi-LAN argued that the plain language of the claims requires that the SDUs be from only one ("a") specified connection. JA1095-1097.

Further, faced with the district court's construction of claim 1 of the '640 patent to require *multiple* connections, Wi-LAN argued that the accused products have multiple "UL connections," as that term is properly construed. JA1078-1088. Specifically, Wi-LAN argued that the district court's construction of "UL

connections" must be corrected to be located on the "uplink" or "between the subscriber unit and the base station," and not a connection between the subscriber radio unit and the end user application processor. *Id.*

Just as the district court denied that Apple waived its right to make new claim construction arguments in its summary judgment papers, the district court denied that Wi-LAN had waived its right to present its claim construction for "UL connection" in its reconsideration papers and considered Wi-LAN's arguments and evidence on the merits. JA4-5. However, the district court ultimately disagreed with Wi-LAN's arguments and evidence and maintained its claim construction. *Id.*

Following the district court's denial of Wi-LAN's Motion for Reconsideration, the district court entered Final Judgment of Non-Infringement. JA1-2. Wi-LAN now appeals these claim construction issues and the non-infringement ruling based on these constructions.

IV. SUMMARY OF ARGUMENT

1. For the '640 patent, the proper construction of "UL connection" is between a subscriber unit and a base station. The district court's construction placing this connection between the subscriber unit and the end user is plain error. The district court did not identify any supporting evidence for its construction. JA27; JA4-5. The patentee defined the "UL [uplink] connection" as existing between the subscriber unit and the base station. JA73 (1:51-52) ("Transmissions

from the subscriber unit to the base station are commonly referred to as 'uplink' transmissions."). The district court's decision to ignore the patentee's chosen definition is plain error.

Further, the district court's construction of "UL connections" significantly contradicts the intrinsic record, including i) the claim language which recites communication only between the subscriber unit and the cellular base station, ii) the specification which describes and illustrates uplink connections between the subscriber unit and the base station, iii) the prosecution history in which the Patent Examiner interpreted "UL connections" as located between the subscriber unit and the base station, and iv) the aim of the patent to optimize the wireless communication link from the subscriber unit to the base station where cellular "uplink" bandwidth is supply-limited.

2. For the '040 patent, the district court's construction of "associated with a specified connection" as requiring the SDUs come from more than one connection is erroneous. [redacted] The district court essentially re-wrote the claims to impose a requirement of more than one "specified connection" and did so based on a faulty reading of the claim language and patent specification. That was error. The patentee is entitled to the full scope of his claims absent a "clear and unmistakable disavowal of claim scope." No such disavowal of claim scope was

found by the district court. Indeed, the invention of the '040 is about packing SDUs from "a" specified connection into a PDU. Moreover, the district court's construction cannot be correct because it i) ignores the plain language of the claims, ii) excludes a preferred embodiment, and iii) violates claim differentiation.

V. ARGUMENT

A. Standards of Review

1. Claim Construction

The Supreme Court recently reaffirmed that claim construction is a question of law. *Teva Pharm. USA, Inc. v. Sandoz, Inc.*, 135 S. Ct. 831, 837 (U.S. 2015). The Supreme Court held that the district court's subsidiary factual determinations may be entitled to deference, but the district court in this case did not make any subsidiary factual findings. Thus, this Court reviews the district court's claim construction de novo. *Id.*; *Schindler Elevator Corp. v. Otis Elevator Co.*, 593 F.3d 1275, 1281 (Fed. Cir. 2010); *Cybor Corp. v. FAS Techs., Inc.*, 138 F.3d 1448, 1454-55 (Fed. Cir. 1998) (*en banc*).

2. Summary Judgment

A summary judgment order can stand only if the record shows "that there is no genuine dispute as to any material fact and [that Defendants are] entitled to judgment as a matter of law." Fed. R. Civ. P. 56(a). In reviewing a summary judgment order, this Court "resolv[es] reasonable factual inferences in favor of the

patentee." *Absolute Software, Inc. v. Stealth Signal, Inc.,* 659 F.3d 1121, 1130 (Fed. Cir. 2011). In other words, Wi-LAN's summary judgment evidence "is to be believed, and all justifiable inferences are to be drawn in [its] favor." *Anderson v. Liberty Lobby, Inc.,* 477 U.S. 242, 255 (1986).

B. This Court Should Reverse Summary Judgment for the '640 Patent Because the Proper Construction of "UL Connection" Is Between the Subscriber Unit and the Base Station, Not Between the Subscriber Unit and an End User, as Construed

The district court has never cited any evidence to support its construction of "UL connection" as being a connection between the subscriber unit and the end user. JA27; JA4-5. The district court's construction is erroneous, as shown below.

1. The Patentee Defined the "Uplink" As Located Between the Subscriber Unit and the Base Station

It is not disputed that the term "uplink" (abbreviated "UL")[7] is between the subscriber unit and the base station, as expressly defined in the '640 patent:

> Transmissions **from the subscriber unit to the base station** are commonly referred to as ***'uplink'*** transmissions.

JA73 (1:51-52).

Therefore, an uplink connection is a connection between the subscriber unit and the base station, not a connection between a subscriber unit and an end user (as construed). The district court departed from the patentee's definition and chose a

[7] As noted earlier, "[t]here is no dispute 'UL' mean 'uplink.'" JA27.

new definition for "uplink" without providing any basis or citing any evidence. JA27. This is plain error. *Martek Biosciences Corp. v. Nutrinova, Inc.*, 579 F.3d 1363, 1380 (Fed. Cir. 2009) ("When a patentee explicitly defines a claim term in the patent specification, the patentee's definition controls.").

2. The Claims Describe Communication Occurring Only Between the Subscriber Unit and Base Station

Claim 1 of the '640 patent is directed entirely to communication between the subscriber unit and the base station. The claim recites registering and "establishing communication between the wireless subscriber radio unit and the base station." JA84 (23:13-14). The claim also describes a two-step process between the subscriber unit and the base station for requesting an uplink bandwidth grant, i.e., an allocation of uplink bandwidth in which to transmit data from the subscriber unit to the base station. JA84 (23:15-27). Finally, the claim teaches allocating the uplink bandwidth grant to one or more UL connections. JA84 (23:27-33).

Thus, the claims focus squarely and solely on the communication between the subscriber unit and the base station and do not claim a link between a subscriber unit and an end user or describe communication on any link to an end user. JA84 (23:7-33). Moreover, the patent optimizes the use of the limited wireless bandwidth assigned to cellular traffic and is not concerned with optimizing the use of bandwidth for communication between the subscriber radio

unit and the end user application processor, which may be wired or wireless. *Supra* Section IV.B.; JA2480 (¶8).

Further, dependent claim 5 of the ’640 patent recites that “UL connections [are] ***established at the wireless subscriber radio unit***.” JA84 (23:45-48). When the subscriber unit has data in one of its queues pending transport to the base station, it establishes a “UL connection” for transporting the data to the base station. JA79 (14:46-57); JA82 (20:6-8) (“connections are established . . . based upon data pending”). In short, the patentee defined the term “uplink” to be between the subscriber unit and the base station, and the patent describes and illustrates uplink connections between the subscriber unit and the base station.

3. The Specification Describes and Illustrates Uplink Connections

The ’640 patent describes and illustrates connections that are established (i.e. originate) at the subscriber radio unit to transport data to a base station over the wireless network, as shown in Figure 10 below.

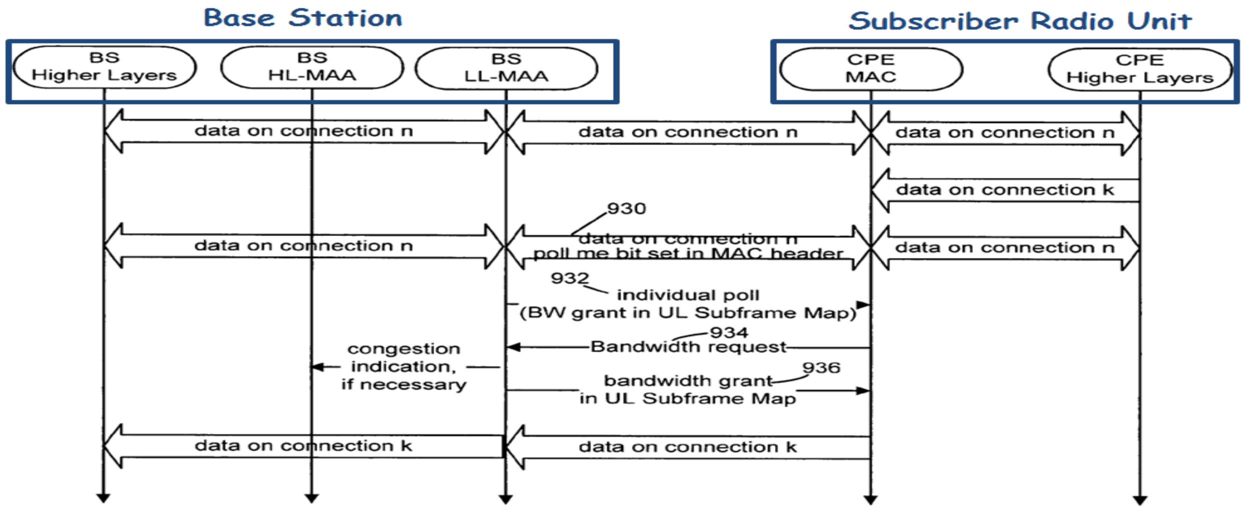

FIGURE 10

JA69 (annotated). As illustrated above, the subscriber radio unit ("CPE") allocates bandwidth first to higher priority "connection n" for uplink data communication. While communicating data to the base station ("BS") for "connection n," the subscriber unit has data needing transmission for lower priority "connection k." The subscriber unit performs a two-step bandwidth request process (930, 932, 934) to obtain additional bandwidth (936), which it allocates to "connection k." *Supra* Section IV.B.1; JA81 (17:57-18:4).

4. The Patent Examiner and Patentee Interpreted "UL Connection" As Between the Subscriber Unit and the Base Station

When attempting to apply prior art to the pending claims, the Patent Examiner construed the term "UL connection" just as Wi-LAN proposes. Specifically, the Patent Examiner read the voice and multimedia data connections in the prior art Basu patent (U.S. Patent No. 6,097,733), which exist between the

subscriber radio unit and the base station as shown below, on the claimed "UL connections." JA1271; JA1384, JA1536, and JA1616-1617. The patentee did not dispute that these connections in Basu are "UL connections," and distinguished Basu on other grounds. JA1405.[8]

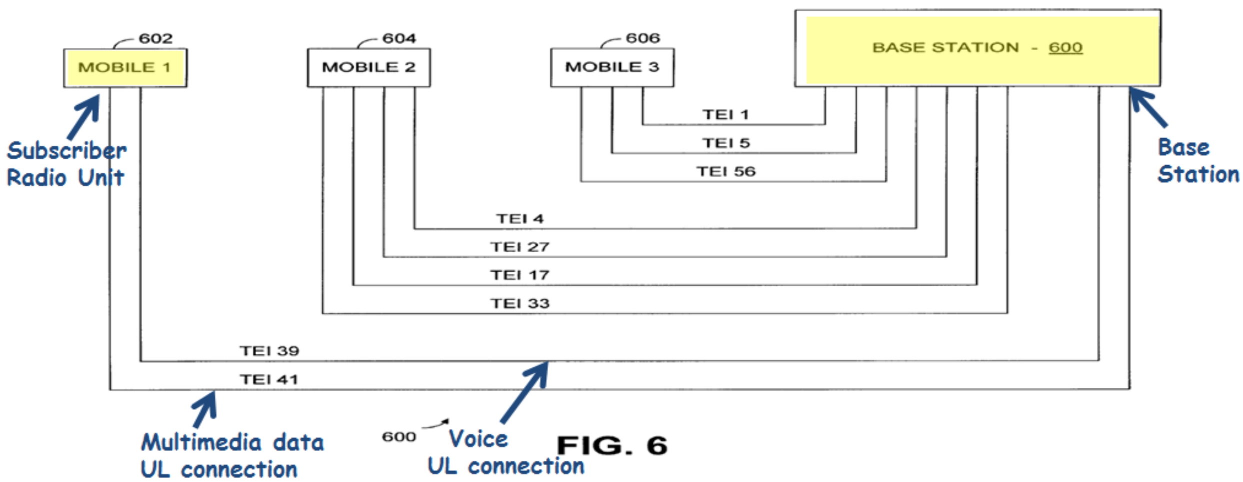

JA1636 (annotated). Thus, the Patent Examiner's and patentee's interpretation that a "UL connection" is between the subscriber unit and the base station further shows that the district court's construction was erroneous. *Cordis Corp. v. Medtronic Ave,* 339 F.3d 1352, 1359-60 (Fed. Cir. 2003) (reversing the district

[8] Additionally, the Patent Examiner's treatment of original claim 18 shows that he interpreted "UL connection" as a distinct and separate link than a connection to an end user. The Patent Examiner found that the connections described in original claim 18 were connections going back to end users, as opposed to the "UL connections" found in other original claims, and thus issued a restriction requirement requiring the patentee to elect claims for continued prosecution. JA1240. Following the restriction requirement, the patentee elected to cancel claim 18 and proceed with prosecution of the "UL connection" claims. JA1257.

court's construction of a term because it was inconsistent with the Patent Examiner's interpretation).

5. Apple Cannot Rebut That the Claimed "UL Connections" Are Between the Subscriber Unit and the Base Station

Apple has never addressed the definition of "uplink" in the '640 patent, which is dispositive that the recited "UL [or uplink] connection" is located between the subscriber unit and the base station. *Martek*, 579 F.3d at 1380 (Fed. Cir. 2009). The crux of Apple's argument appears to be that because the allocation of bandwidth to an uplink connection is *for* an end user, it must necessarily mean the "UL connection" is the connection from the subscriber radio unit to the end user. JA2515-2516. Not so. Apple relies on the following passages:

> CPEs 110 request bandwidth allocations from their respective base stations 106 based upon the type and quality of services *requested by the customers* served by the CPEs. Different broadband services have different bandwidth and latency requirements.

JA73 (2:16-33) (italics added).

> The CPE is responsible for distributing the allocated uplink bandwidth in a manner that accommodates the services provided by the CPE.

JA74 (4:34-36).

Bandwidth is allocated on the uplink and downlink connections (i.e. the "links" between the subscriber radio unit and the base station) where the "bottleneck" in the transmission path occurs, i.e., where bandwidth is precious.

Multiple subscriber radio units are competing to send their data on uplink connections using an allocation of the supply-limited uplink bandwidth.[9] *Id.* The base station, which grants bandwidth in the first instance, must grant enough uplink bandwidth to support the bandwidth needs of the subscriber unit so that it can meet the priority or quality of service requirements of its end user and control data. *Id.* Once the base station has granted bandwidth to a subscriber unit, it is for the subscriber unit to prioritize its various pending data transmissions such that the higher priority data is generally transmitted on a "UL connection" before lower priority data. JA82 (19:28-31).

In addition, Apple has argued that claim 6 of the '640 patent, a claim which is not at issue on appeal, shows that a "UL connection" is between the end user and the subscriber radio unit. JA2492. But, significantly, "UL connection" is not used in claim 6, which instead uses the generic term "connection." Moreover, the particular claim language Apple identifies, "traffic . . . received on a plurality of connections" from unasserted claim 6 (JA84 (24:9-10)) is not present in asserted claim 1.

In short, Apple cannot (and has not even attempted to) rebut the patentee's express definition of "UL [uplink] connection" to be between the subscriber unit

[9] In contrast, data transmitted between an end user and the subscriber unit can move over a wired connector or over unlicensed bandwidth that can, and does, transmit data at significantly higher speeds than the cellular network. JA2480 (¶8).

and the cellular base station. The district court's construction is erroneous and Apple cannot show otherwise.

C. This Court Should Reverse Summary Judgment of Non-Infringement of the '040 Patent Because "Associated with a Specified Connection" Requires Only One Specified Connection

The district court construed the language each SDU "associated with a specified connection" to effectively add the word "***different***" before the words "specified connection." [redacted]

[redacted]

[redacted]

[redacted]

[redacted] Moreover, the claims do not say that each SDU is associated with a "***different***" specified connection, and, in fact, confirm the opposite:

> packing and fragmenting the ***SDUs associated with the specified connection*** into the payload area of the PDU based on the current length of the payload area;

JA56 (21:27-29).

In its Order denying Wi-LAN's Motion for Reconsideration, the district court pointed to an embodiment in Figures 3 and 7 and several specification passages as allegedly teaching that the term "associated with a specified connection" requires multiple connections. JA5. However, the embodiments pointed out by the district court merely show that the subscriber unit node may accept SDUs from more than one end user, not that each SDU that is packed and fragmented *must* (or even should) be associated with a ***different*** specified connection.

Moreover, the error in the district court's reasoning is further highlighted by the fact that it contradicts the claim language and the prosecution history, violates claim differentiation, and excludes a preferred embodiment.

[10] The only evidence the district court cited is the claim language to show that multiple SDUs are at issue, which is not in dispute. JA11. The district court cited no evidence which suggests, let alone requires, that each SDU must be associated with a *different* specified connection.

1. The District Court's Construction Improperly Narrows the Plain Language of the Claims Without Finding Disclaimer

The '040 patent consistently and repeatedly claims a node for packing and fragmenting multiple SDUs associated with **one** specified connection into a PDU, as shown in claim 1 below.[11] This is underscored by the claim language, which refers back to "a specified connection" as **"the** specified connection," confirming only one "specified connection" is required.

> 1. A node for a communications system that packs and fragments variable-length service data units (SDU) for mapping into variable length protocol data units (PDU), each SDU being **associated with a specified connection**, the node comprising:
>
> a communications processor configured to pack and fragment **SDUs associated with a specified connection** into a PDU, including
>
> allocate bandwidth for **the specified connection**, based on the priority of the connection,
>
> establish a length for the PDU based on the bandwidth allocated to **the specified connection** in a current frame,
>
> ***

JA55 (19:29-40). Independent claim 16 is also directed to packing and fragmenting multiple SDUs associated with one "specified connection" into a PDU. Here, the language is always "**the** specified connection," showing that one connection, not multiple, is required.

[11] The claim language of independent claim 14 is not recited here because it is similar to that of claim 1.

> 16. A method of formatting protocol data units (PDUs) from incoming variable-sized service data units (SDUs) for transmission of data carried by the PDUs over a communication channel shared by one or more nodes, comprising:
>
> provisioning a protocol data unit (PDU), including a header and a payload area, wherein the length of the PDU is established in conjunction with the bandwidth amount allocated to **<u>the</u> specified connection** in a current frame, the bandwidth amount being established frame-by-frame based on one or more communication parameters associated with **<u>the</u> specified connection,** including the priority of **<u>the</u> specified connection**, and general system parameters;
>
> packing and fragmenting **<u>the</u> SDUs associated with <u>the</u> specified connection** into the payload area of the PDU based on the current length of the payload area;
>
> ***

JA56 (21:13-29). The scope of protection afforded by the plain language of claims may be limited only upon a "clear and unmistakable" disclaimer of claim scope:

> It is likewise not enough that the only embodiments, or all of the embodiments, contain a particular limitation. We do not read limitations from the specification into claims; we do not redefine words. Only the patentee can do that. To constitute disclaimer, there must be a ***clear and unmistakable*** disclaimer . . . . The patentee is free to choose a broad term and expect to obtain ***the full scope of its plain and ordinary meaning unless the patentee explicitly redefines the term or disavows its full scope.***

Thorner v. Sony, 669 F.3d 1362, 1366-67 (Fed. Cir. 2012). [redacted]

Even if the district court's description of how the invention allegedly processes SDUs had merit (it does not, as described below), the district court cannot rewrite the claim language, unless there has been a "clear and unmistakable" disclaimer. *Thorner*, 669 F.3d at 1366-67 (Fed. Cir. 2012); *Rhine v. Casio, Inc.*, 183 F.3d 1342, 1345 (Fed. Cir. 1999) (finding the district court's construction at odds with the clear language of the claim and written description when it construed "at least *one* light source" to require at least two light sources).

Finally, the claim language "***each SDU being*** associated with a specified connection" exists only in the preamble of independent claim 1 of the '040 patent. In contrast, the body of claim 1 confirms that only one specified connection is required. Similarly, independent claim 16 recites "packing and fragmenting **the** SDUs associated with **the** specified connection," JA56 (21:27-28), and claim 14 recites "packs and fragments **SDUs** associated with **a** specified connection." JA20 (20:52-54). Neither claim 14 nor claim 16 recite the claim language "each SDU being associated with a specified connection," further showing each SDU that is packed and fragmented need not be associated with a ***different*** "specified connection."

2. The District Court's Construction Contradicts the Prosecution History

During the prosecution history of the parent patent to the '040 patent, the patentee made clear that the "specified connection" is the connection whose associated SDUs are packed and fragmented into a PDU:

> In addition, there is no teaching in *Van Grinsven* to establish the length for the PDU based on bandwidth allocated to the specified connection (**the connection whose associated SDUs are being packed and fragmented into the PDU**) in the current frame as set forth in the independent claims.

JA1917. This intrinsic evidence further shows that the SDUs are associated with one specified connection, in contradiction to the district court's requirement that each SDU must be associated with a *different* specified connection.[12]

3. The District Court's Construction Violates Claim Differentiation Principles

The district court fails to recognize the claim differentiation problem that arises from dependent claim 7, which refers back to "a specified connection" as "**the** connection," and requires that the header identify one specified connection (not connections plural):

[12] The '040 patent is a continuation of U.S. Patent No. 8,009,667 (the "'667 parent patent"). JA29. Evidence from the '667 parent patent prosecution history is intrinsic evidence to the '040 patent. *Goldenberg v. Cytogen*, 373 F.3d 1158, 1167 (Fed. Cir. 2004) (treating parent patent file history as intrinsic evidence).

> 7. A node as claimed in claim 6, wherein the communications processor is further configured to **identify** **<u>the connection</u>** in the header.

JA55 (20:4-6). Courts "must not interpret an independent claim in a way that is inconsistent with a claim which depends from it." *Wright Med. Tech., Inc. v. Osteonics Corp.,* 122 F.3d 1440, 1445 (Fed. Cir. 1997).

Even Apple does not dispute the claim differentiation problem. JA2512 (lines 8-13) (not disputing the claim differentiation problem and making other arguments). Further, the preferred embodiment claimed in dependent claim 7 is illustrated in Figure 8 of the patent where "the connection" is identified by a single "connection identifier" in the header of the PDU, as further discussed below.

4. The District Court's Construction Excludes a Preferred Embodiment

The district court's construction of "associated with a specified connection" excludes a preferred embodiment — packing multiple SDUs from the <u>same</u> connection, not multiple connections, to save bandwidth:

> Packing ***multiple SDUs from the <u>same connection</u>*** behind a single MAC header ***saves bandwidth*** over attaching a separate MAC header to each.

JA2120. Consistent with this teaching, the '040 patent teaches that the specified connection is identified by a single 16-bit connection identifier (<u>not</u> connections plural) (JA2487 (¶28)), as shown below:

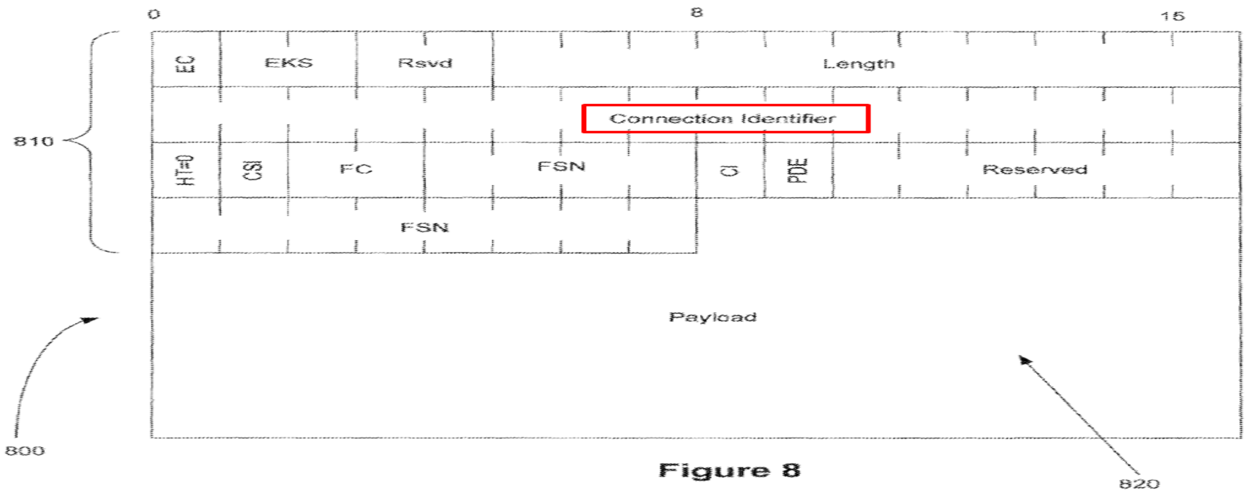

Figure 8

JA39 (Fig. 8). The fact that the district court's construction excludes a preferred embodiment that packs multiple SDUs from the "same connection" further shows that the district court's construction is erroneous. *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1583 (Fed. Cir. 1996) (a claim interpretation that excludes a preferred embodiment is "rarely, if ever, correct").

5. Apple Does Not Deny that "Associated With A Specified Connection" Requires Only One Specified Connection

Tellingly, Apple does not defend the district court's reasoning, thereby implicitly admitting the fallacy in the district court's conclusion that the SDUs must each be from a ***different*** "specified connection." In response to Wi-LAN's reconsideration motion, Apple pointed to entirely separate claim language, in bold below, in an erroneous attempt to salvage a multiple "specified connection" requirement:

> allocate bandwidth for the specified connection, **based on the priority of the connection**

> establish a length for the PDU based on the bandwidth allocated to the specified connection in a current frame

JA55 (19:37-40). Apple made the following argument:

> [Wi-LAN] argue[s] that SDUs from the ***same*** connection are packed into a single PDU. To the extent this embodiment is illustrated in the '040 Patent, the '040 Patent's allocation and prioritization scheme still requires that the length of each PDU be determined as a result of the priority-based allocation of bandwidth to each of a plurality of [specified] connections.

JA2512. At bottom, Apple argues that the claim language "based on the priority of the connection" must mean "based on the priority of the connection ***in relation to multiple specified connection***s." But Apple's reasoning is flawed.

After bandwidth is allocated by the subscriber unit to meet the priority requirements of the SDUs for a "specified connection," the node establishes a length for the PDU into which the SDUs are packed and fragmented. JA46 (2:3-5), JA55 (19:39-40). But there is nothing in the claims that requires that the allocation of bandwidth to the "specified connection" *must* be based on the priority of the connection in relation to another "specified connection." JA2490 (¶38). The inventors recited their invention broadly and are entitled to the full scope of the claim invention absent a clear and unmistakable disclaimer. *Thorner*, 669 F.3d at 1366 ("The patentee is free to choose a broad term and expect to obtain the full

scope of its plain and ordinary meaning unless the patentee explicitly redefines the term or disavows its full scope.”).

Apple’s reasoning is also erroneous because it would exclude an embodiment in which a “specified connection,” carrying end user data as construed (JA24), is prioritized in relation to “control data” for a “control connection” [redacted] The ’040 patent teaches a “control connection” and “control data” at the subscriber radio unit having a QoS priority in relation to end user data from a specified connection, and vice versa. JA2218 (“Bandwidth Allocation and Request Mechanisms[:] Note that at registration every SS [subscriber unit] is assigned *two dedicated CIDs* [connection identifiers] for the purpose of sending and receiving *control messages*. Two connections are used to allow *differentiated levels of QoS* to be applied to different connections carrying MAC management traffic.”); JA2144 (“Definitions: Connection Identifier (CID): A unidirectional, MAC-layer address that identifies a connection to equivalent peers in the SS and BS MAC. A CID maps to a SFID [Service Flow Identifier], *which defines the QoS parameters* . . .”).[13]

[13] The quoted passages appear in the provisional application, which is *incorporated by reference in its entirety into the ’040 patent specification*. JA46 (1:7-11).

In short, Apple abandons the district court's reasoning in its response to Wi-LAN's reconsideration motion and, through its own faulty reasoning, attempts to limit the claims to require multiple "specified connections" through a distinct claim term ("based on the priority of the connection") not addressed by the district court. The district court's construction of "associated with a specified connection" is erroneous and Apple cannot show otherwise.

D. Wi-LAN Can Show Infringement Under the Proper Constructions, and the District Court's Grant of Summary Judgment Should, Therefore, Be Reversed and Remanded

1. The '640 Patent

Under the proper construction of "UL connection," i.e. a connection between a subscriber unit and the base station, Wi-LAN can show infringement. [redacted]

[redacted]

At the very least, Wi-LAN presents a fact question that makes summary judgment

inappropriate. Fed. R. Civ. P. 56(a); *Absolute Software*, 659 F.3d at 1130 (Fed. Cir. 2011). Accordingly, summary judgment of the ’640 patent claims should be reversed and the case remanded to the district court.

2. The ’040 Patent

[redacted]

[redacted]

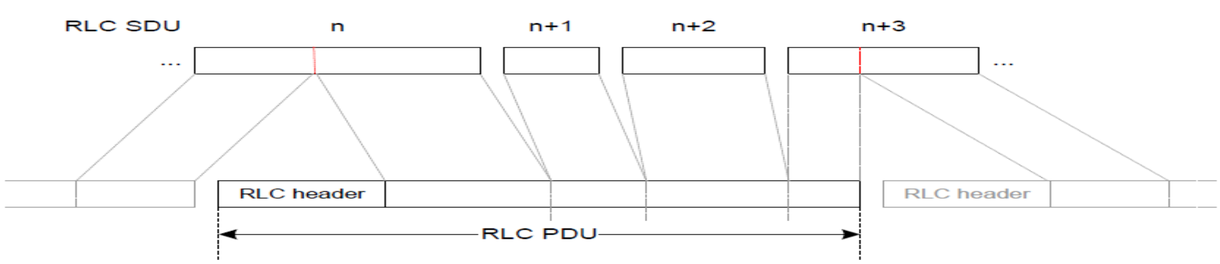

Figure 6.2.2-1: RLC PDU Structure

[redacted]

[redacted] At the very least, Wi-LAN presents a fact question making summary judgment inappropriate. Fed. R. Civ. P. 56(a); *Absolute Software,* 659 F.3d at 1130 (Fed. Cir. 2011). Accordingly, summary judgment of non-infringement of the ’040 patent claims should be reversed and the case remanded to the district court.

VI. CONCLUSION

Wi-LAN respectfully requests, based on the foregoing, that this Court reverse the district court’s grant of summary judgment of all claims and remand with instructions regarding the proper constructions of “UL connections” as connections between the subscriber unit and the base station and “associated with a specified connection” as requiring only one specified connection.

Dated: March 13, 2015

Respectfully submitted,

/s/ Robert A. Cote

Robert A. Cote
Kevin Schubert
McKOOL SMITH P.C.
One Bryant Park, 47th Floor
New York, New York 10036
Tel: (212) 402-9400
Fax: (212) 402-9444
rcote@mckoolsmith.com
kschubert@mckoolsmith.com

Dirk D. Thomas
McKOOL SMITH P.C.
1999 K Street, Suite 600
Washington, DC 20006
Tel: (202) 370-8302
Fax: (202) 370-8344
DThomas@McKoolSmith.com

Attorneys for Plaintiffs-Appellants, Wi-LAN USA, Inc., and Wi-LAN Inc.

ADDENDUM

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

SAN DIEGO DIVISION

WI-LAN USA, INC. and WI-LAN, INC.

Plaintiff,

v.

APPLE INC.

Defendant.

Case No. 3:13-cv-798-DMS (BLM)

JOINT STIPULATED FINAL JUDGMENT

Judge: Hon. Dana M. Sabraw
Magistrate Judge: Hon. Barbara L. Major

Complaint Filed: December 6, 2012

STIPULATED FINAL JUDGMENT

Pursuant to this Court's Orders Construing Patent Claims ("Markman") (Doc. No. 98), Granting Plaintiffs' Motion for Clarification and Reconsideration of the Markman Order (Doc. No. 123), Granting Defendant's Motion for Summary Judgment of Non-Infringement as to all Claims (Doc. No. 278), and Denying Plaintiffs' Motion for Reconsideration of the Summary Judgment Order (Doc. No. 299), the Court hereby enters final judgment in favor of Defendant Apple Inc. ("Apple") on Plaintiffs Wi-LAN USA, Inc. and Wi-LAN, Inc.'s claims for infringement of U.S. Patent No. 8,311,040 and U.S. Patent No. 8,315,640 (and Apple's corresponding claims for declaratory relief of non-infringement).

It is further **ORDERED, ADJUDGED and DECREED** that all remaining claims, counterclaims, and affirmative defenses be, and hereby are, dismissed without prejudice as MOOT. The Court, however, will allow Apple to re-assert any and all such claims, counterclaims, and affirmative defenses at a later point in time if this Court's order on non-infringement is vacated on appeal.

It is further **ORDERED, ADJUDGED and DECREED** that Apple is the prevailing party in this matter and is entitled to costs pursuant to F.R.C.P. 54(d)(1).

The Clerk of the Court is directed to close this case.

So ORDERED and SIGNED this 8th day of January, 2015.

HON. DANA M. SABRAW
UNITED STATES DISTRICT JUDGE

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

WI-LAN USA, INC. and WI-LAN, INC., Plaintiffs, vs. APPLE INC., Defendant. ――― AND ALL RELATED COUNTERCLAIMS.	CASE NO. 13cv0798 DMS (BLM) **ORDER DENYING WI-LAN USA, INC. AND WI-LAN INC.'S MOTION FOR RECONSIDERATION**

On September 30, 2014, this Court issued an Order granting Apple's motion for summary judgment of noninfringement of the Patents in Suit. On October 28, 2014, the Wi-Lan parties ("Wi-Lan") filed the present motion. They seek reconsideration of the Court's decision that Apple does not infringe the '040 Patent and claims 1 and 2 of the '640 Patent.[1] Apple filed an opposition to the motion, and Wi-Lan filed a reply.

"Reconsideration is appropriate if the district court (1) is presented with newly discovered evidence, (2) committed clear error or the initial decision was manifestly unjust, or (3) if there is an intervening change in controlling law." *School Dist. No. 1J,*

[1] Wi-Lan does not seek reconsideration of the Court's decision that Apple does not infringe claims 6 and 7 of the '640 Patent.

Multnomah County, Oregon v. ACandS, Inc., 5 F.3d 1255, 1263 (9th Cir. 1993). Here, Wi-Lan argues the Court committed clear error in its constructions of "specified connection" in the '040 Patent and "UL connections" in the '640 Patent. Wi-Lan also asserts the Court committed clear error in interpreting the '040 Patent to require more than one "specified connection."

With respect to the Court's construction of the term "specified connection" as "the communications link between a node and a specific end user" and "UL connections" as "an uplink connection between the wireless subscriber radio unit and its users," Wi-Lan has not shown the Court committed clear error. Notably, the present motion is the first time Wi-Lan has challenged the Court's constructions of these terms, despite filing a motion for reconsideration of other terms and the passage of more than ten months since the Court issued its rulings. The impetus for Wi-Lan's motion appears to be the Court's finding of noninfringement, which although adverse to Wi-Lan, does not demonstrate the Court's claim construction is clearly erroneous. The evidence Wi-Lan relies on to support its argument of clear error is not newly discovered. It has been available to Wi-Lan since the outset of this case, as have the arguments it raises in the present motion.

Nevertheless, the Court has considered Wi-Lan's arguments and its newly produced evidence, and finds they do not demonstrate the Court's claim construction is clearly erroneous. With respect to the '040 Patent, although the Patentee amended the claims in the parent application to use the term "specified connection" instead of "user connection," that amendment does not demonstrate the Patentee intended to change the connection that was being described. (*See* Decl. of Dirk Thomas in Supp. of Mot. ("Thomas Decl."), Ex. 5 at 1164.) Furthermore, the evidence Wi-Lan relies on refers to a "connection," not a "specified connection" as that term is used in the claims. (*Id.* at 1301; Thomas Decl., Ex. 6 at 1531.) With respect to the '640 Patent, the Court has reviewed the newly provided evidence, and finds it does not demonstrate the

/ / /

Court's construction of "UL connection" is clearly erroneous. Therefore, the Court declines to reconsider its constructions of these terms.

Wi-Lan has also failed to show that the Court's interpretation of the '040 Patent as requiring more than one "specified connection" is clearly erroneous. Figures 3 and 7 support the Court's finding that more than one "specified connection" is required. So do many other portions of the specification. (*See*, *e.g.*, '040 Patent at 4:40-41, 50-51) (describing nodes as serving "multiple connections for users" and transmitting information "to the appropriate connection for transmission to the appropriate user"); (*id.* at 6:19-20) (stating connection interface is "coupled to a plurality of user connections"). Accordingly, the Court declines to reconsider its finding that the asserted claims of the '040 Patent require more than one "specified connection."

For these reasons, Wi-Lan's motion for reconsideration is denied. A telephonic status conference shall be held on **December 18, 2014**, at **2:30 p.m.** to discuss the further progress of this case. Counsel for Wi-Lan shall organize and initiate the conference call to the Court.

IT IS SO ORDERED.

DATED: December 12, 2014

HON. DANA M. SABRAW
United States District Judge

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

REDACTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

WI-LAN USA, INC. and WI-LAN, INC., Plaintiffs, vs. APPLE INC., Defendant.	CASE NO. 13cv0798 DMS (BLM) **ORDER GRANTING PLAINTIFFS' MOTION FOR CLARIFICATION AND RECONSIDERATION** **[Docket No. 108]**

On December 23, 2013, this Court issued its Order Construing Patent Claims. In that Order, the Court construed the term "node" in the '040 Patent as "a module between a base station and an end user that directs transmission of data over a communications link." The Court also construed the terms "wireless subscriber radio unit/wireless communication radio unit" in the '640 Patent as a "module that receives UL bandwidth from a base station, and allocates the bandwidth across its user connections." Another term at issue was "packing sub-header," which the Court construed as "a header located in a PDU payload."

On January 21, 2014, Wi-Lan filed the present motion for clarification and reconsideration of the Court's Claim Construction Order, seeking clarification of the Court's constructions of "node" and "wireless subscriber radio unit/wireless communication radio unit," and reconsideration of the Court's construction of "packing sub-header."[1] On the former, the parties disagree whether the Court's

[1] The Court notes that Wi-Lan substituted new counsel after the Court's *Markman* Order and before filing the present motion.

constructions of the terms "node" and "wireless subscriber radio unit/wireless communication radio unit" preclude Wi-Lan from arguing that these aspects of the invention can be "a component part of a user device." (Mem. of P. & A. in Supp. of Mot. at 3.) Wi-Lan argues the Court's construction does not preclude it from making this argument, while Apple argues Wi-Lan is so precluded. On the latter, Wi-Lan requests the Court reconsider its construction of the term "packing sub-header." The motion came on for hearing on April 2, 2014. Dirk Thomas and Seth Hasenour appeared for Wi-Lan and Mark Scarsi and Miguel Ruiz appeared for Apple. After thoroughly reviewing the parties' briefs and the record on file herein, and hearing oral argument from counsel, the Court grants Wi-Lan's motion.

Clarification of the Court's constructions of "node" and "wireless subscriber radio unit/wireless communication radio unit" is appropriate and necessary. Obviously, the parties disagree about the effect of the Court's claim construction, and the Court needs to resolve that disagreement. Furthermore, "[d]istrict courts may engage in a rolling claim construction, in which the court revisits and alters its interpretation of the claim terms as its understanding of the technology evolves." *Jack Guttman, Inc. v. Kopykake Enterprises, Inc.*, 302 F.3d 1352, 1361 (Fed. Cir. 2002) (citing *Sofamor Danek Group, Inc. v. DePuy-Motech, Inc.*, 74 F.3d 1216, 1221 (Fed. Cir. 1996)).

Wi-Lan argues the Court's claim construction does not preclude it from arguing that the "node" and "wireless subscriber radio unit/wireless communication radio unit" are components of a user device, such as a cellular phone or tablet computer. Wi-Lan asserts there is nothing inherent in the Court's constructions that prevents it from making this argument, and the intrinsic evidence supports Wi-Lan's position. Apple argues Wi-Lan made this argument during the claim construction process, the Court rejected it then and should reject it now.

The Court agrees with Apple that Wi-Lan made this same argument during the claim construction process, *i.e.*, that the "node" and "wireless subscriber radio unit/wireless subscriber communication unit" could be a component of a cellular phone or tablet computer. The Court also agrees that it rejected that argument when it adopted Apple's proposed constructions of these terms. However, Wi-Lan has presented evidence that was not presented during the claim construction process. Specifically, Wi-Lan points to portions of the prosecution history for the '640 Patent that

/ / /

reflect the Examiner was reading prior art concerning a mobile unit on the "wireless subscriber radio unit" in the '640 Patent.

The Court has reviewed this evidence, and finds it supports Wi-Lan's position that the "wireless subscriber radio unit/wireless communication radio unit" could be a component of a cellular phone or tablet computer. Throughout the prosecution history, the Examiner compared the "wireless subscriber radio unit/wireless communication radio unit" to prior art containing a mobile device. (*See* Ex. E in Supp. of Mot. at 18-22, 37-38, 40, 42-43, 65-70, 99, 103, 105-06, 117-20, 122, 124-27.)

Apple did not address this evidence in its brief, but did address the evidence at oral argument. Specifically, Apple argued that the Examiner focused on a prior art reference, namely United States Patent Number 6,016,311 to Gilbert, that contained an intermediary device. However, as Wi-Lan pointed out at oral argument, the Gilbert patent referred to in the prosecution history is not the '311 Patent. Rather, it is Gilbert's United States Patent Number 5,297,144, (*see id.* at 65), which does not require an intermediary device. Apple's argument does not refute Wi-Lan's interpretation of the prosecution history. Rather, the evidence supports Wi-Lan's position that it should be allowed to argue that the "wireless subscriber radio unit/wireless communication radio unit" could be a component part of a user device. The Court so finds, and thus clarifies its construction of these terms.[2]

Turning to the term "packing sub-header," the parties' arguments are of the same nature. Wi-Lan makes the same argument it made during the claim construction process, but relies on different evidence, and Apple argues the Court has already rejected Wi-Lan's argument once and should do so again.

Here, again, the Court agrees with Wi-Lan, and finds reconsideration of the Court's claim construction is in order. *See School Dist. No. 1J v. AcandS, Inc.*, 5 F.3d 1255, 1263 (9th Cir. 1993) (quoting *United States v. U.S. Gypsum Co.*, 333 U.S. 364, 395 (1948)) (stating reconsideration is appropriate "when 'the reviewing court on the entire record is left with the definite and firm

---

[2] Throughout the claim construction proceedings, Wi-Lan has treated the terms "wireless subscriber radio unit/wireless communication radio unit" in the '640 Patent the same as the term "node" in the '040 Patent. It continues that approach in the present motion, asserting that the term "node" should be interpreted consistently with the terms "wireless subscriber radio unit/wireless communication radio unit." The Court agrees these terms should be construed consistently, and thus clarifies the term "node" is interpreted in accordance with the terms "wireless subscriber radio unit/wireless communication radio unit."

conviction that a mistake has been committed.'") Although the argument is the same, Wi-Lan points to portions of the specification that were not called out during the claim construction process. Specifically, it relies on Figure 8 of the '040 Patent, combined with the dependent claims, which show the "packing sub-header" in the PDU header but not in the PDU payload.[3]

Apple did not address this evidence in its brief, but did address it at oral argument. There, Apple argued Figure 8 did not include a packing sub-header. However, the Court disagrees with that argument. Dependent claims 4, 5, 11 and 12 state a packing sub-header may comprise a fragment sequence number ("FSN") or a fragmentation control ("FC") field. Figure 8 clearly shows both a FSN and FC field in the PDU header, not in the PDU payload. In light of this evidence, the Court grants Wi-Lan's motion for reconsideration of this issue, and amends the claim construction to remove the limitation that the "packing sub-header" is located in the PDU payload.

IT IS SO ORDERED.

DATED: April 10, 2014

HON. DANA M. SABRAW
United States District Judge

[3] At oral argument, Wi-Lan raised additional arguments and evidence that were not included in its briefs. The Court finds consideration of those arguments and evidence is unnecessary to the present motion. Thus, Wi-Lan need not file that additional paperwork with the Court, nor does Apple need to file a response.

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

WI-LAN USA, INC. and WI-LAN, INC., Plaintiffs, vs. APPLE INC., Defendant.	CASE NO. 13cv0798 DMS (BLM)
AND RELATED COUNTERCLAIMS.	**ORDER CONSTRUING PATENT CLAIMS**

This matter came before the Court for a claim construction hearing on December 12, 2013. David B. Weaver, Steven R. Borgman and Andrea Houston appeared and argued on behalf of Plaintiffs Wi-Lan USA, Inc. and Wi-Lan Inc. ("Wi-Lan"). Mark Scarsi appeared and argued on behalf of Apple Inc. After a thorough review of the parties' claim construction briefs and all other material submitted in connection with the hearing, the Court issues the following order construing the disputed terms of the patents at issue in this case.

I.

BACKGROUND

On December 6, 2012, Wi-Lan filed the present Complaint against Apple alleging claims of infringement of United States Patent Numbers 8,311,040 ("the '040 Patent") and 8,315,640 ("the '640 Patent"). The Complaint was filed in the United States District Court for the Southern District of Florida. In response, Apple filed an Answer and Counterclaims for noninfringement, invalidity,

unclean hands, waiver and estoppel, and exceptional case under 35 U.S.C. § 285. On Apple's motion to transfer venue, the case was transferred to this Court on April 2, 2013.

II.

DISCUSSION

Claim construction is an issue of law, *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 372 (1996), and it begins "with the words of the claim." *Nystrom v. TREX Co., Inc.*, 424 F.3d 1136, 1142 (Fed. Cir. 2005) (citing *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996)). Generally, those words are "given their ordinary and customary meaning." *Id.* (citing *Vitronics*, 90 F.3d at 1582). This "'is the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention.'" *Id.* (quoting *Phillips v. AWH Corp.*, 415 F.3d 1303, 1313 (Fed. Cir. 2005)). "The person of ordinary skill in the art views the claim term in the light of the entire intrinsic record." *Id.* Accordingly, the Court must read the claims "'in view of the specification, of which they are a part.'" *Id.* (quoting *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 979 (Fed. Cir. 1995)). In addition, "'the prosecution history can often inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution, making the claim scope narrower than it would otherwise be.'" *Id.* (quoting *Phillips*, 415 F.3d at 1318).

Here, there are two patents at issue: the '040 Patent and the '640 Patent. The '040 Patent is entitled, "Packing Source Data Packets Into Transporting Packets With Fragmentation." Wi-Lan alleges Apple is infringing claims 1, 2, 4, 5, 6, 7, 8, 9, 10, 13, 14, 16 and 22 of the '040 Patent. The '640 Patent is entitled, "Methods and Systems for Transmission of Multiple Modulated Signals Over Wireless Networks." Wi-Lan alleges Apple is infringing claims 1, 2, 4, 5, 6 and 7 of the '640 Patent. Of the asserted claims, there are six terms that require construction, which the Court discusses below.

A. The '040 Patent

All of the disputed terms in the '040 Patent are found in claim 1, which recites:[1]

> 1. A **node** for a communications system that packs and fragments variable-length service data units (SDU) for mapping into variable length protocol data units (PDU), each SDU being associated with a **specified connection**, the node comprising:

[1] Disputed terms are in bold and underlined.

> a communications processor configured to pack and fragment SDUs associated with a specified connection into a PDU, including
>
> allocate **bandwidth** for the specified connection, based on the priority of the connection,
>
> establish a length for the PDU based on the bandwidth allocated to the specified connection in a current frame,
>
> pack a first SDU into a payload area of the PDU,
>
> determine whether a second SDU is larger than a remaining payload area of the PDU,
>
> if the second SDU is not larger than the remaining payload area of the PDU, map the second SDU to the remaining payload area of the PDU, and
>
> if the second SDU is larger than the remaining payload area of the PDU, fragment the second SDU into at least two fragments and map the first fragment to the remaining payload area of the PDU, and
>
> include **packing sub-headers** in the PDU to allow determination of the length of the SDUs and the lengths of the fragments that are mapped to the PDU.

1. "Node"

The first term at issue is "node." Wi-Lan asserts this term should be construed as "a fixed, portable or mobile wireless unit,"[2] while Apple argues the term should be construed as "a module between a base station and an end user, that directs transmission of data over a communications link."

Both sides rely on the specification to support their proposed constructions. However, Wi-Lan's proposed construction finds little support therein. Rather, the specification provides greater support to Apple's proposed construction. For instance, the specification states each node serves "multiple connections *for users*." ('040 Patent at 4:40-41) (emphasis added). It then goes on to describe "users" as "a service network such as a LAN, WAN, Intranet, Ring Network or other type of network; or they may be a single user such as a work station." (*Id.* at 41-44.)[3] In describing the communications process, the specification states: "Information is received by the base station 12 from the data source, is prepared for and transmitted across a data link to a node 16, and is then directed to

/ / /

[2] Wi-Lan offers this same construction for different terms in the '640 Patent, namely "wireless subscriber radio unit" and "wireless communication radio unit."

[3] Wi-Lan argues this language indicates the node may be a single user, but the language belies that argument. The specification is describing various types of *users*, not types of nodes.

the appropriate connection for transmission to the appropriate *user*." (*Id.* at 4: 47-51) (emphasis added).

This idea of the node as an intermediary between the base station and end users is found throughout the specification. (*See id.* at 4:57-62) (stating node may "discard any data not pertinent to *the users* on its connections.") (emphasis added); (*id.* at 6:11-12) (stating communications processor "converts the signal into the SDUs that *the users* had transmitted to the node 16.); (*id.* at 11:14-30) ("The information intended for each node 16 contains information to be distributed to the *end users* or services served by the connections of that node; ... The connection data 720 includes the information to be transmitted to *the users* or services as well as control information the node uses to identify to which of its connections each packet of information should be directed. Thus, the nodes can ensure that each of the packets of information that it received is directed to the appropriate connection to reach the intended *end user* or service.") (emphasis added); (*id.* at 14:38-42) (describing transformation of data from PDU format "back to the various SDU packet formats that were originally received by the nodes *from the users*.") (emphasis added). (*See also id.*, Figs. 3, 7.) Thus, the Court adopts Apple's proposed construction of the term "node" as "a module between a base station and an end user that directs transmission of data over a communications link."

2. <u>"Specified Connection"</u>

The second term at issue in the '040 Patent is "specified connection." This term appears in claim 1, as set out above. Wi-Lan argues this term should be construed as "specified service." Apple asserts it should be construed as "the communications link between a node module and a specific end user."

As with the term "node," the Court adopts Apple's proposed construction of "specified connection." Wi-Lan argues the specification uses "services" and "connections" interchangeably, but the specification refutes that argument. For instance, the specification speaks of end users or services as being "served by the connections" of the node. (*Id.* at 11:16.) This description of the system is inconsistent with Wi-Lan's proposed construction. Indeed, it is unclear how a service could be served by itself.

/ / /

Apple's proposed construction is supported by the specification, and it is consistent with Apple's proposed construction of "node." As Apple points out, the specification describes "multiple connections" *between* end users and the "node." (*Id.* at 4:40-41, 47-51, 57-62; 6:16-20; 11:14-17, 17-24.) The Figures of the '040 Patent also support this construction. (*Id.*, Figs. 3, 7.) Accordingly, the Court construes "specified connection" as "the communications link between a node and a specific end user."

3. "Packing Sub-Header"

The third term at issue in the '040 Patent is "packing sub-header." Wi-Lan argues "packing sub-header" should be construed as "additional header information," while Apple asserts "packing sub-header" should be construed as "a header located in a PDU payload." The parties agree a "packing sub-header" has a particular structure (located within the PDU) and function (to indicate the length of the SDU packed inside the PDU). Their dispute centers on whether the "packing sub-header" must be included within the PDU payload area.

Wi-Lan argues the "packing sub-header" is not limited to the PDU payload area because it can be contiguous with or separate from the SDU it represents. The parties do not appear to dispute this issue, and it is supported by the specification. (*See id.* at 2:9-14.) However, it does not support Wi-Lan's argument that the "packing sub-header" need not be included in the PDU payload area.

Apple asserts the only descriptions and drawings of the "packing sub-header" reflect it is included in the PDU payload area, therefore the term should be construed accordingly. The specification supports Apple's proposed construction. (*See id.* at 18:15-61; Fig. 14.) Therefore, the Court adopts Apple's proposed construction and construes "packing sub-header" as "a header located in a PDU payload."

4. "Bandwidth"

The final term at issue in the '040 Patent is "bandwidth." Wi-Lan argues this term should be construed as "data-carrying capacity." Apple asserts it should be construed as "an amount of data that can be transmitted in a particular time period."[4]

/ / /

[4] The parties agree this term should be construed consistently across both patents.

In support of its proposed construction, Wi-Lan relies on the specification of the '040 Patent. It asserts the patentee expressly defined "bandwidth" as "data-carrying capacity," (*see id.* at 1:22-23), therefore the Court should adopt that construction. However, the Court disagrees. To act as his own lexicographer, the patentee "must clearly redefine a claim term 'so as to put a reasonable competitor or one reasonably skilled in the art on notice that the patentee intended to so redefine that claim term.'" *Elekta Instruments S.A. v. O.U.R. Scientific Int'l, Inc.*, 214 F.3d 1302, 1307 (Fed. Cir. 2000) (quoting *Process Control Corp. v. HydReclaim Corp.*, 190 F.3d 1350, 1357 (Fed. Cir. 1999)). In this case, it is not clear that the patentee intended to redefine bandwidth, as Wi-Lan suggests.

In contrast to Wi-Lan, who relies primarily on the specification of the '040 Patent, Apple relies more on the specification of the '640 Patent to support its proposed construction of "bandwidth." However, the specification does not support Apple's position. Apple also relies on an order construing "bandwidth" in other Wi-Lan patents related to the '640 Patent to support its position, but that order is not binding on this Court, and the Court is not persuaded by that court's reasoning in coming to its construction of "bandwidth."

Although the patentee did not expressly define the term "bandwidth" as data-carrying capacity, the Court finds that construction is more consistent with the inventions described in both the '040 and '640 Patents. Throughout the Patents, "bandwidth" is described as a commodity of the system, something that is allocated to various links, ('040 Patent at 9:32-39; '640 Patent at abstract), and distributed by the base station. ('040 Patent at 13:44-48; '640 Patent at 1:20-24.) Contrary to Apple's suggestion, it does not describe a unit of speed. Rather, as Apple acknowledges, it describes the capacity of the device or system to send information. Thus, based on the Court's review of the intrinsic evidence, the Court construes "bandwidth" as "data-carrying capacity."

B. The '640 Patent

Turning to the '640 Patent, all of the disputed claim terms are found in claim 1, which recites:[5]

> 1. A method for requesting **bandwidth** on demand in a wireless communication system, wherein the wireless communication system includes a **wireless subscriber radio unit**, the method comprising:

[5] Disputed terms are in bold and underlined.

> registering the **wireless communication radio unit** with a base station in the wireless communication system and establishing communication between the wireless subscriber radio unit and the base station;
>
> transmitting from the wireless subscriber radio unit which is registered with the base station, an explicit message to the base station requesting to be provided an allocation of uplink (UL) bandwidth in which to transmit a bandwidth request;
>
> receiving at the wireless subscriber radio unit the allocation of UL bandwidth in which to transmit a bandwidth request;
>
> transmitting the bandwidth request within the allocation of UL bandwidth, the bandwidth request specifying a requested UL bandwidth allocation; and
>
> receiving an UL bandwidth grant for the wireless subscriber radio unit in response to the bandwidth request;
>
> wherein the wireless subscriber radio unit maintains a plurality of queues, each queue for data pertaining to one or more **UL connections** with similar QoS and wherein the wireless subscriber radio unit allocates the UL bandwidth grant to the one or more UL connections based on QoS priority.

1. "Wireless Subscriber Radio Unit/Wireless Communication Radio Unit"

The first terms at issue in the '640 Patent are "wireless subscriber radio unit/wireless communication radio unit." These terms raise disputes similar to those raised by the term "node" in the '040 Patent. Indeed, Wi-Lan offers the same proposed construction for these terms as they did for the term "node," namely "a fixed, portable or mobile wireless unit." Apple argues the terms should be construed as "customer premises equipment that receives UL bandwidth from a base station, and allocates the bandwidth across its user connections."

For the reasons set out above in the discussion of "node," the Court agrees with Apple that the "wireless subscriber radio unit/wireless communication radio unit" sits in a similar position as the "node" in the '040 Patent: Both are intermediaries between the base station and end users.

The primary dispute, then, is whether the "wireless subscriber radio units/wireless communication radio units" are equivalent to customer premises equipment ("CPE").[6] Although the specification refers repeatedly to CPEs as part of the invention, it is not clear that CPEs are interchangeable with "wireless subscriber radio units/wireless communication radio units." For

[6] At oral argument, Apple conceded it was not necessary to construe the "wireless subscriber radio units" and "wireless communication radio units" as CPEs. (Hearing Tr. at 81-82.) Therefore, this issue may now be moot. However, to the extent it is not, the Court addresses it below.

instance, the specification describes "fixed subscriber stations *or* Customer Premises Equipment (CPE)." ('640 Patent at 1:62-65) (emphasis added). Furthermore, the specification refers to "subscriber stations," not wireless radio units. Finally, at oral argument, Apple stated the patentee amended the claims to recite "wireless subscriber radio units" and "wireless subscriber communication units" rather than CPEs. (Hearing Tr. at 56.) Under these circumstances, the Court declines to limit "wireless subscriber radio units" and "wireless communications radio units" to CPEs. Rather, the Court construes these terms as a "module that receives UL bandwidth from a base station, and allocates the bandwidth across its user connections."[7]

2. "UL Connections"

The final term at issue is "UL connections." There is no dispute "UL" means uplink. The dispute is how to construe "connections." Wi-Lan argues, consistent with its proposed construction of "specified connection" in the '040 Patent, that "connection" should be construed as "service," while Apple asserts it should be construed as a connection between the CPE and its users.

As it did with the '040 Patent, Wi-Lan argues the specification of the '640 Patent equates "connections" with "services," therefore the Court should adopt its proposed construction. However, the Court again disagrees with that argument. The specification describes connections and services, but nowhere does it equate the two.

Apple's proposed construction is more consistent with the Court's understanding of the invention, which describes connections between end users and wireless subscriber radio units.[8] Therefore, the Court construes this term as "an uplink connection between the wireless subscriber radio unit and its users."

/ / /

/ / /

[7] In the briefing, there appeared to be a dispute as to whether the "wireless subscriber radio units" and "wireless communication radio units," as well as the "nodes" in the '040 Patent, were fixed, portable or mobile. At oral argument, Apple appeared to concede that these modules "could be either fixed or portable." (*Id.* at 57-58.) Given that concession, the Court declines to address that issue further.

[8] Apple uses the term CPE instead of "wireless subscriber radio unit." As discussed above, the Court declines to impose that limitation on the claims.

III.

CONCLUSION

For the reasons stated above, the disputed terms are interpreted as set forth in this Order.

IT IS SO ORDERED.

DATED: December 23, 2013

HON. DANA M. SABRAW
United States District Judge

US008311040B2

(12) **United States Patent**
Stanwood et al.

(10) **Patent No.: US 8,311,040 B2**
(45) **Date of Patent: *Nov. 13, 2012**

(54) **PACKING SOURCE DATA PACKETS INTO TRANSPORTING PACKETS WITH FRAGMENTATION**

(75) Inventors: **Kenneth L. Stanwood**, Carlsbad, CA (US); **Stanley Wang**, San Diego, CA (US); **Robert M. Johnson**, San Diego, CA (US)

(73) Assignee: **Wi-LAN, Inc.**, Ottawa (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 218 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/886,314**

(22) Filed: **Sep. 20, 2010**

(65) **Prior Publication Data**

US 2011/0033048 A1 Feb. 10, 2011

Related U.S. Application Data

(63) Continuation of application No. 10/053,179, filed on Jan. 15, 2002, now Pat. No. 8,009,667.

(60) Provisional application No. 60/262,005, filed on Jan. 16, 2001.

(51) **Int. Cl.**
H04L 12/56 (2006.01)
H04J 3/24 (2006.01)

(52) **U.S. Cl.** **370/389**; 370/471; 370/473

(58) **Field of Classification Search** .... None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,949,404	A	4/1976	Fletcher et al.
4,495,619	A	1/1985	Acampora
4,975,906	A	12/1990	Takiyasu et al.
5,130,983	A	7/1992	Heffner, III
5,297,144	A	3/1994	Gilbert et al.
5,402,422	A	3/1995	Liu et al.
5,404,374	A	4/1995	Mullins et al.
5,420,851	A	5/1995	Seshadri et al.
5,444,698	A	8/1995	Kito
5,457,687	A	10/1995	Newman
5,511,082	A	4/1996	How et al.
5,517,503	A	5/1996	Hess

(Continued)

FOREIGN PATENT DOCUMENTS

EP 0507384 3/1992

(Continued)

OTHER PUBLICATIONS

Cable Television Laboratories, Inc., "Data-Over-Cable Service Interface Specifications DOCSIS 1.1, Radio Frequency Interface Specification", 1999-2005.

(Continued)

Primary Examiner — Gregory Sefcheck
(74) *Attorney, Agent, or Firm* — Procopio, Cory, Hargreaves & Savitch LLP

(57) **ABSTRACT**

A communication system and method are disclosed for transmitting packets of information in at least one first format over a communications link that utilizes packets of information in a second format. In certain embodiments, the packets of information in a first format are converted to packets of information in the second format prior to transmission via the communications link by packing and fragmenting the information in the first format in a coordinated manner. Embodiments may also utilize packing subheaders and fragmentation control bits in the packing and fragmentation processes.

22 Claims, 14 Drawing Sheets

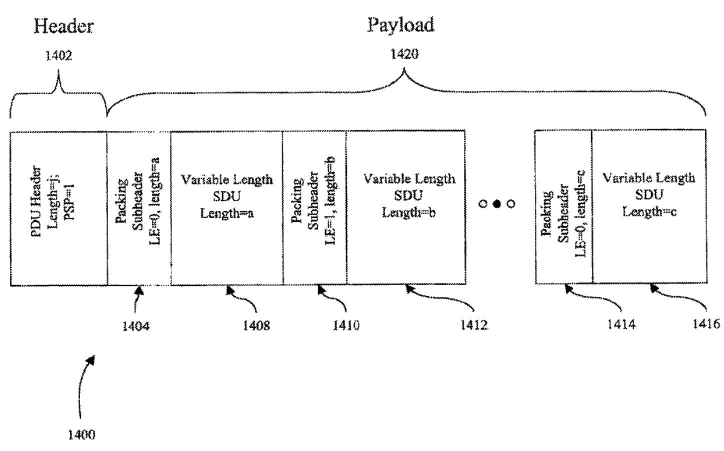

U.S. PATENT DOCUMENTS

5,615,212	A	3/1997	Ruszczyk et al.	
5,638,371	A	6/1997	Raychaudhuri et al.	
5,638,374	A	6/1997	Heath	
5,666,358	A	9/1997	Paratore et al.	
5,675,573	A	10/1997	Karol et al.	
5,684,791	A	11/1997	Raychaudhuri et al.	
5,751,708	A	5/1998	Eng et al.	
5,768,254	A	6/1998	Papadopoulos et al.	
5,774,461	A	6/1998	Hyden et al.	
5,818,842	A	10/1998	Burwell et al.	
5,828,695	A	10/1998	Webb	
5,844,906	A	12/1998	Khelghatti et al.	
5,859,619	A	1/1999	Wu et al.	
5,859,848	A	1/1999	Miura et al.	
5,890,055	A	3/1999	Chu et al.	
5,909,469	A	6/1999	Frodigh et al.	
5,912,920	A	6/1999	Marchok et al.	
5,918,171	A	6/1999	Funke et al.	
5,946,313	A	8/1999	Allan et al.	
5,963,543	A	10/1999	Rostoker et al.	
5,970,062	A	10/1999	Bauchot	
6,006,069	A	12/1999	Langston	
6,011,786	A	1/2000	Dent	
6,016,311	A	1/2000	Gilbert et al.	
6,016,313	A	1/2000	Foster, Jr. et al.	
6,016,401	A	1/2000	Rostoker et al.	
6,026,133	A	2/2000	Sokoler	
6,031,827	A	2/2000	Rikkinen et al.	
6,038,455	A	3/2000	Gardner et al.	
6,094,421	A	7/2000	Scott	
6,112,080	A	8/2000	Anderson et al.	
6,128,293	A	10/2000	Pfeffer	
6,147,989	A	11/2000	Esaki et al.	
6,151,318	A	11/2000	Woodward et al.	
6,167,031	A	12/2000	Olofsson et al.	
6,188,680	B1	2/2001	Ohki	
6,189,041	B1	2/2001	Cox et al.	
6,201,811	B1	3/2001	Larsson et al.	
6,226,277	B1 *	5/2001	Chuah	370/328
6,233,231	B1	5/2001	Felix et al.	
6,249,530	B1 *	6/2001	Blanco et al.	370/468
6,339,585	B1	1/2002	Hulyalkar et al.	
6,351,499	B1	2/2002	Paulraj et al.	
6,359,934	B1	3/2002	Yoshida	
6,366,961	B1	4/2002	Subbiah et al.	
6,370,112	B1	4/2002	Voelker	
6,400,755	B1	6/2002	Harris et al.	
6,452,927	B1	9/2002	Rich	
6,452,964	B1	9/2002	Yoshida	
6,463,107	B1	10/2002	Lindoff et al.	
6,469,991	B1	10/2002	Chuah	
6,493,342	B1	12/2002	Breslow et al.	
6,496,142	B1	12/2002	Iinuma	
6,549,759	B2	4/2003	Arviv et al.	
6,567,383	B1	5/2003	Bohnke	
6,570,849	B1 *	5/2003	Skemer et al.	370/230.1
6,608,841	B1	8/2003	Koodli	
6,631,130	B1	10/2003	Roy et al.	
6,665,292	B1	12/2003	Bradshaw et al.	
6,665,296	B1	12/2003	Sturza et al.	
6,674,731	B1	1/2004	Bradshaw et al.	
6,683,866	B1	1/2004	Stanwood et al.	
6,697,364	B1	2/2004	Kekki et al.	
6,711,176	B1	3/2004	Pezeshki-Esfahani	
6,721,286	B1	4/2004	Williams et al.	
6,741,614	B1	5/2004	Porter et al.	
6,747,977	B1	6/2004	Smith et al.	
6,760,599	B1	7/2004	Uhlik	
6,771,660	B1	8/2004	Bourlas et al.	
6,785,252	B1	8/2004	Zimmerman et al.	
6,804,251	B1	10/2004	Limb et al.	
6,839,823	B1 *	1/2005	See et al.	711/171
6,862,622	B2	3/2005	Jorgensen	
6,891,855	B2 *	5/2005	Bruckman	370/468
6,907,048	B1	6/2005	Treadaway et al.	
6,917,984	B1 *	7/2005	Tan	709/238
6,918,034	B1	7/2005	Sengodan et al.	
6,925,068	B1	8/2005	Stanwood et al.	
6,931,030	B1	8/2005	Dogan	
6,940,823	B1	9/2005	Park et al.	
6,944,672	B2	9/2005	Crow et al.	
6,954,796	B2 *	10/2005	Tan et al.	709/231
6,963,751	B1	11/2005	Kordsmeyer et al.	
6,967,928	B1 *	11/2005	Kikuta	370/252
6,968,554	B1 *	11/2005	MacDonald et al.	719/316
6,970,941	B1	11/2005	Caronni et al.	
7,006,525	B1	2/2006	Jha	
7,006,526	B1	2/2006	Biederman	
7,020,715	B2	3/2006	Venkataraman et al.	
7,050,407	B1	5/2006	Frazer et al.	
7,061,935	B1	6/2006	Roy et al.	
7,068,639	B1	6/2006	Varma et al.	
7,085,291	B2	8/2006	Zhang et al.	
7,197,022	B2	3/2007	Stanwood et al.	
7,197,025	B2 *	3/2007	Chuah	370/338
7,260,405	B2	8/2007	Kim et al.	
7,321,569	B2	1/2008	Takagi et al.	
7,428,247	B2 *	9/2008	Bunn et al.	370/474
7,443,814	B2	10/2008	Agarwal et al.	
7,570,645	B2	8/2009	Agarwal	
7,583,623	B2	9/2009	Zimmerman et al.	
7,693,186	B2 *	4/2010	Bunn et al.	370/474
2001/0015985	A1	8/2001	Van Grinsven et al.	
2001/0048681	A1	12/2001	Bilic et al.	
2002/0016852	A1	2/2002	Nishihara	
2002/0089984	A1	7/2002	Jiang	
2002/0126677	A1	9/2002	Hathaway et al.	
2003/0169722	A1	9/2003	Petrus et al.	
2004/0132459	A1	7/2004	Varma et al.	
2005/0100120	A1	5/2005	Barton et al.	
2005/0220110	A1	10/2005	Agarwal	
2006/0062250	A1	3/2006	Payne	
2006/0262808	A1 *	11/2006	Lin et al.	370/466

FOREIGN PATENT DOCUMENTS

EP	0622924		11/1994
EP	0720405		7/1996
EP	0845916		6/1998
EP	0891060		1/1999
EP	0944200		9/1999
JP	11266231	A	9/1999
WO	WO9222162		12/1992
WO	WO 9713388		4/1997
WO	WO 9810568		3/1998
WO	99/12303	A1	3/1999
WO	99/18685	A1	4/1999
WO	WO 9938343		7/1999
WO	WO 9939532		8/1999
WO	WO 0001188		1/2000

OTHER PUBLICATIONS

Dutton et al., "Asynchronous Transfer Mode (ATM) Technical Overview," 2.sup.nd Edition, Prentice Hall Oct. 1995, Chap. 3, pp. 21-25.

International Search Report for PCT/US00/29687 dated Mar. 28, 2001, 3 pages.

International Preliminary Examination Report for PCT/US00/29687 dated Dec. 18, 2001, 8 pages.

Office Action issued in U.S. Appl. No. 09/430,379 on Jan. 30, 2003, 10 pages.

Notice of Allowance for U.S. Appl. No. 09/430,379 issued on Jul. 28, 2003, 3 pages.

Sampei et al., "Adaptive Modulation/TDMA Scheme for Personal Multi-Media Communication System," Proceedings of the Global Telecommunications Conference, U.S., New York, IEEE, Nov. 28, 1994, pp. 989-993.

Ue et al., "Symbol Rate and Modulation Level Controlled Adaptive Modulation/TDMA/TDD for Personal Communication Systems, Proceedings of the Vehicular Technology Conference, U.S.," New York, IEEE, vol. Conf. 45, Jul. 25, 1995, pp. 306-310.

Torrance et al., "Upper Bound Performance of Adaptive Modulation in a Slow Rayleigh Fading Channel," IEEE Electronics Letters vol. 32, p. 718-719, Apr. 1996.

Pons et al., "Bit Error Rate Characterisation and Modelling for GSM," IEEE 1998, pp. 3722-3727.

Bender et al., “CDMA/HDR: A Bandwidth Efficient High-Speed Wireless Data Service for Nomadic Users,” Communications Magazine, IEEE, vol. 38, No. 7, Jul. 2000, pp. 70-77.
Jain, P., “On the Impact of Channel and Channel Quality Estimation on Adaptive Modulation,” Dec. 2002, 92 pages.
Andersen, et al., “Prediction of Future Fading Based on Past Measurements, Vehicular Technology Conference,” VTC 99, vol. 1, p. 151-155.
Narlikar et al., “Designing Multihop Wireless Backhand Networks with Delay Guarantees,” Bell labs, 2005 12 pages.
U.S. District Court, Northern District of California, “Amended Complaint for Declaratory Judgment Demand for Jury Trial” filed Sep. 30, 2008 in Case No. 5:08-cv-4555, 20 pages.
Lin et al., “Error Control Coding, Fundamentals and Applications”, Prentice-Hall Computer Applications in Electrical Engineering Series., 1993, pp. 315-349.
Lee, L.H. Charles, “Convolutional Coding, Fundamentals and Applications”, Artech House, Inc., 1997, p. 11-51.
Redl, et al., “An Introduction to GSM”, Artech House, Inc., 1995, pp. 84-85 and 95.
Shannon, C.E., “A Mathematical Theory of Communication”, Bell System Technical Journal, pp. 379-423 (Part 1), 623-656 (Part 2), Jul. 1948.
Ulm et al., “Data-Over-Cable Interface Specifications, Radio Frequency Interface Specification”, Hewlett Packard Interim Specification, Doc. Control No. SP-RFII01-970321, published Mar. 21, 1997 by MCNS Holdings, L.P. Section 6, pp. 43-85.
Wolf et al., “On the Weight Distribution of Linear Block Codes Formed From Convolutional Codes”, IEEE, IEEE Transactions on Communications, vol. 44:Sep. 9, 1996, 4 pages.
Papadopoulos et al., “Reduction of Mixed Co-channel Interference in Microcellular STDD Systems”, Vehicular Technology Conference, 1995 IEEE 45.sup.th, vol. 2, pp. 759-763.
Sater et al., “Media Access Control Layer Proposal for the 802.16.1 Air Interface Specification”, IEEE 802.16 Broadband Wireless Access Working Group, 'Online! Jul. 7, 2000, XP002210334 retrieved from the Internet: <URL:http://wirelessman.org/tgl/mac/contrib/802161mc-00 .sub.1321rl. pdf> retrieved Aug. 20, 2002, pp. 69-80.
International Search Report for PCT/US01/43582 dated Aug. 30, 2002, 3 pages.
International Search Report for PCT/US00/20577 dated Apr. 5, 2001, 5 pages.

* cited by examiner

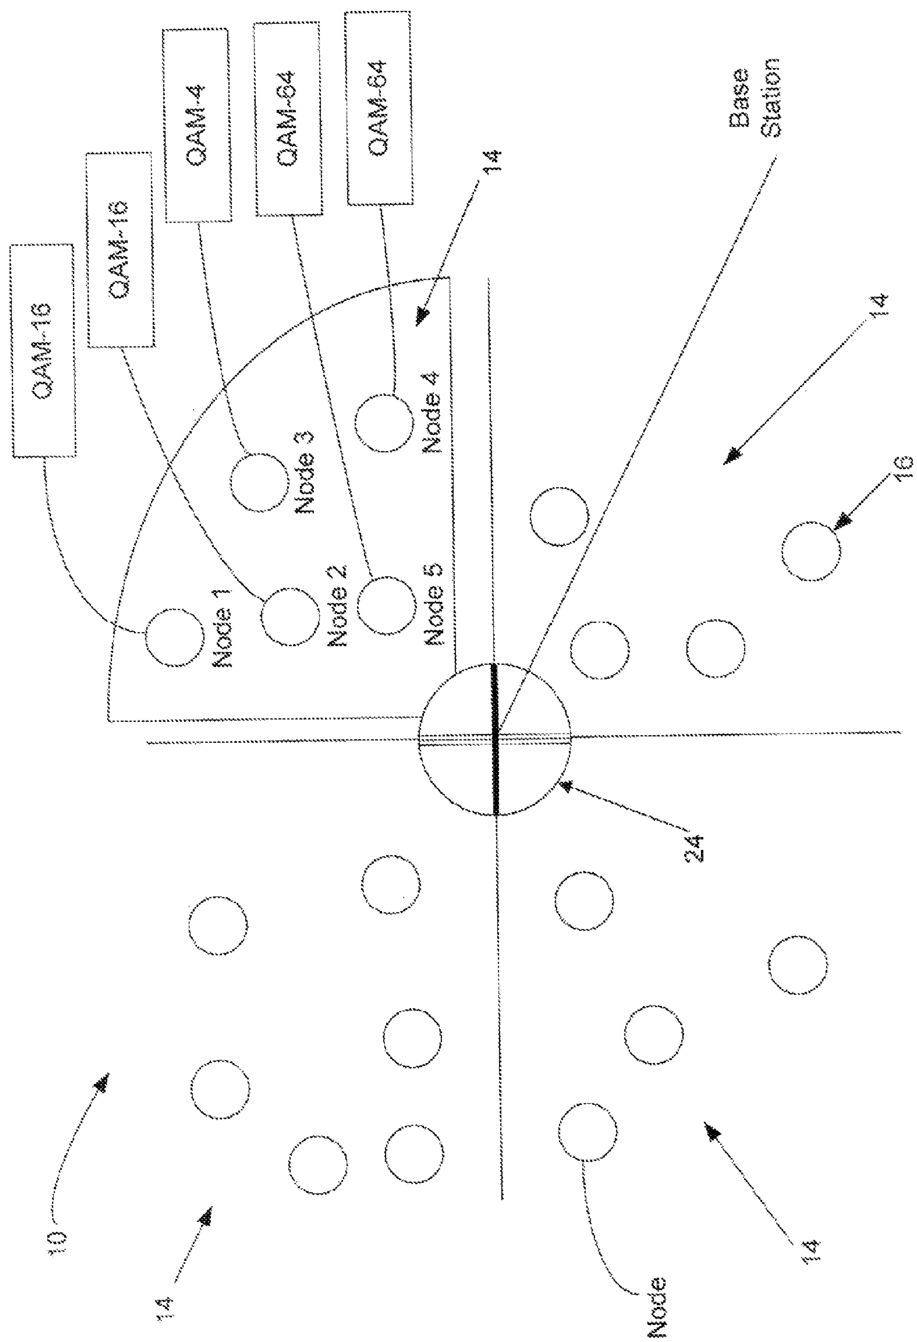
QAM-16
QAM-16
QAM-4
QAM-64
QAM-64
Base Station
14
Node 1
Node 2
Node 3
Node 4
Node 5
14
16
24
10
14
Node
14

Figure 1

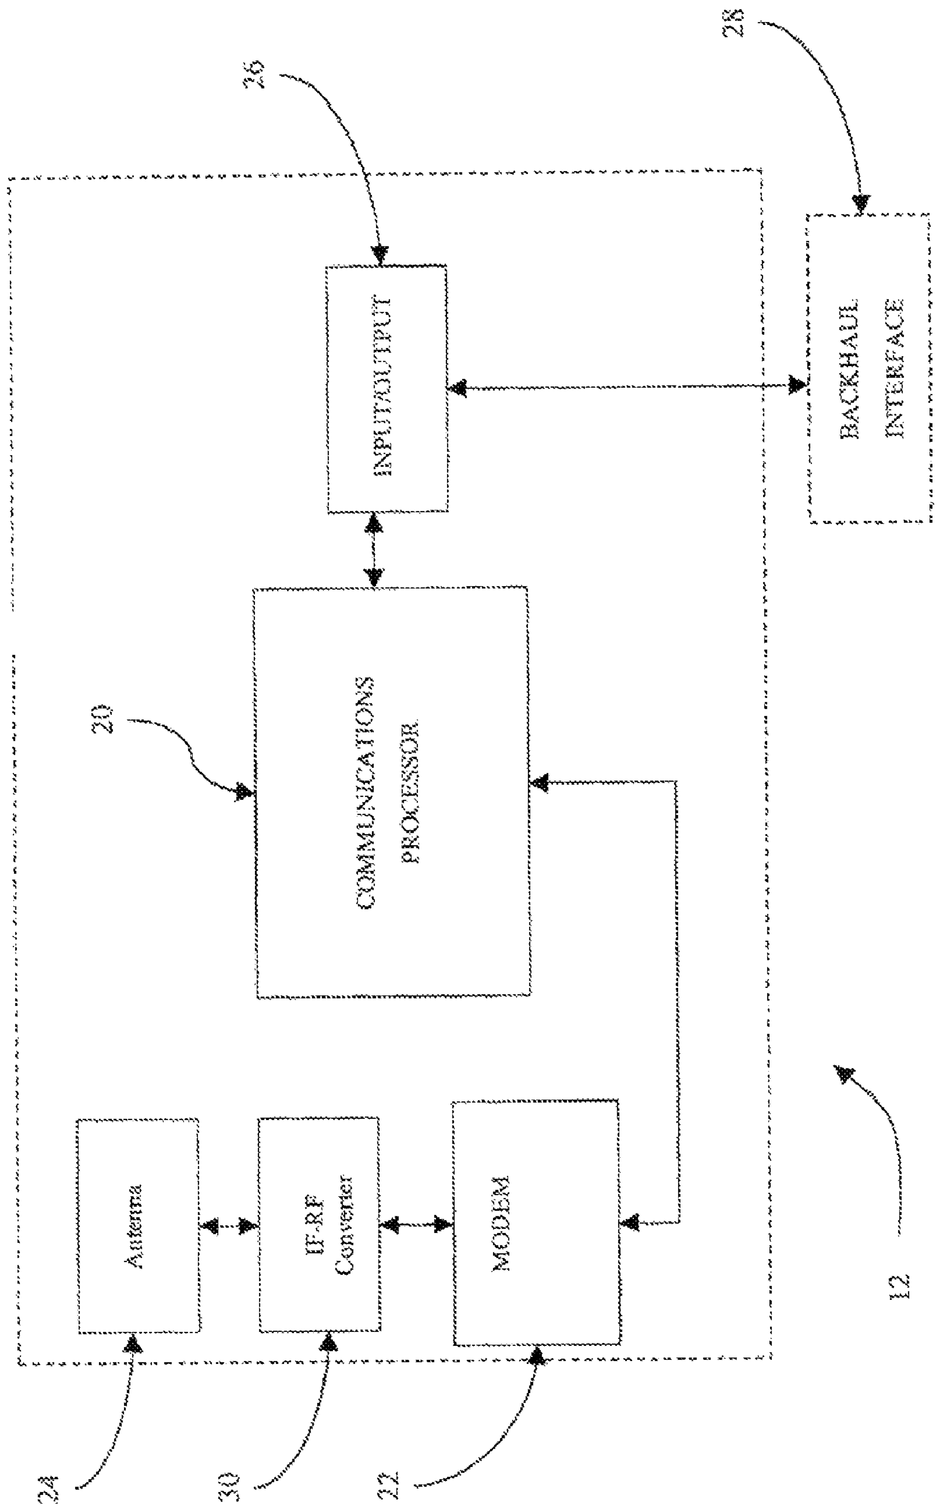
26
28
INPUT/OUTPUT
BACKHAUL INTERFACE
20
COMMUNICATIONS PROCESSOR
Antenna
IF-RF Converter
MODEM
12
24
30
22

Figure 2

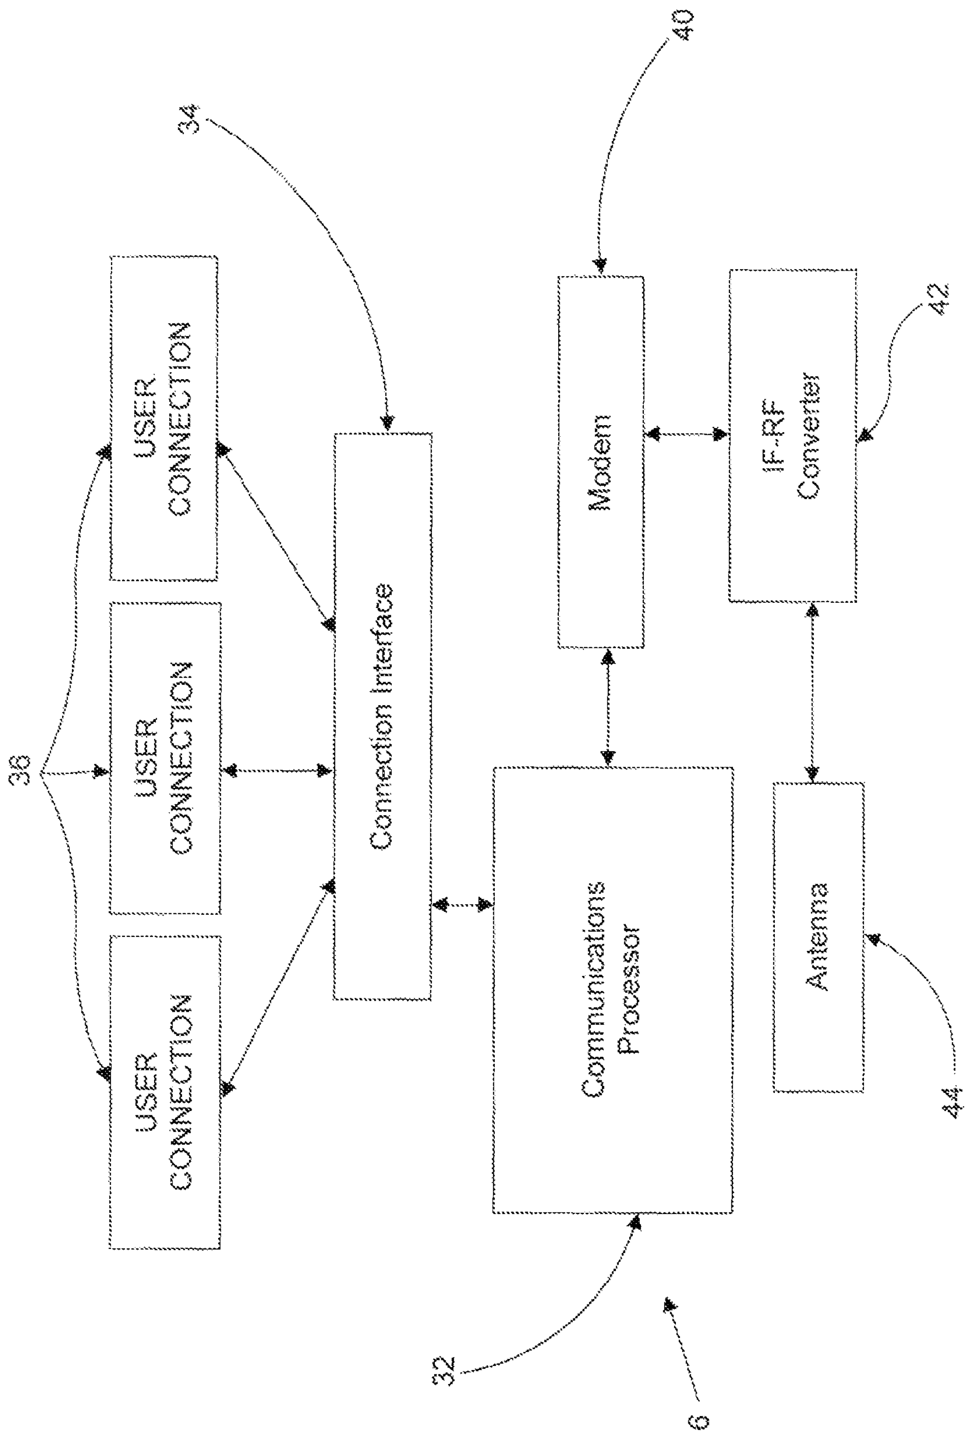
USER CONNECTION
USER CONNECTION
USER CONNECTION
Connection Interface
Modem
IF-RF Converter
Communications Processor
Antenna
16
32
34
36
40
42
44

Figure 3

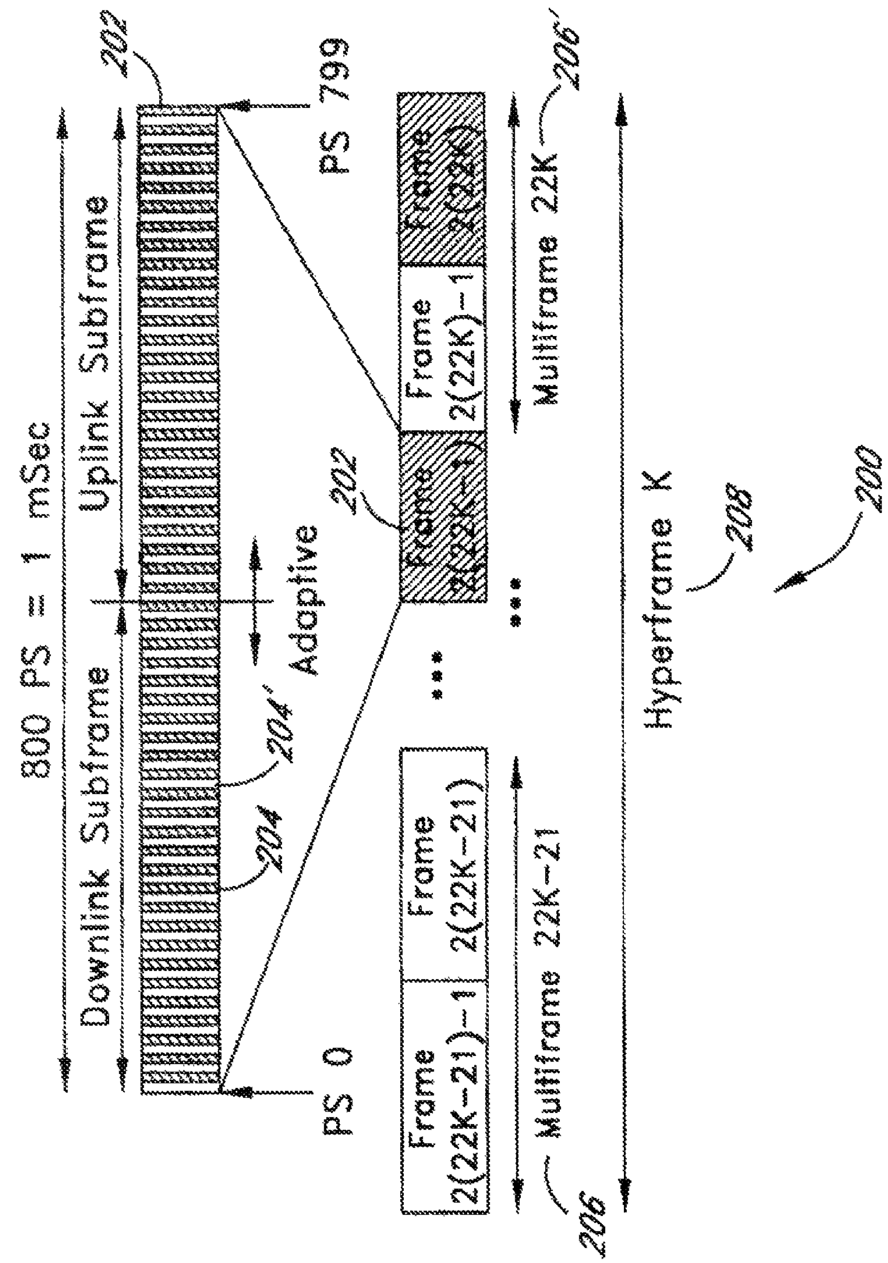
800 PS = 1 mSec
Downlink Subframe
Uplink Subframe
Adaptive
PS 0
PS 799
202
204
204'
Frame 2(22K−21)−1
Frame 2(22K−21)
Frame 2(22K−1)
Frame 2(22K)−1
Frame 2(22K)
Multiframe 22K−21
Multiframe 22K
206
206'
Hyperframe K
208
200

Figure 4

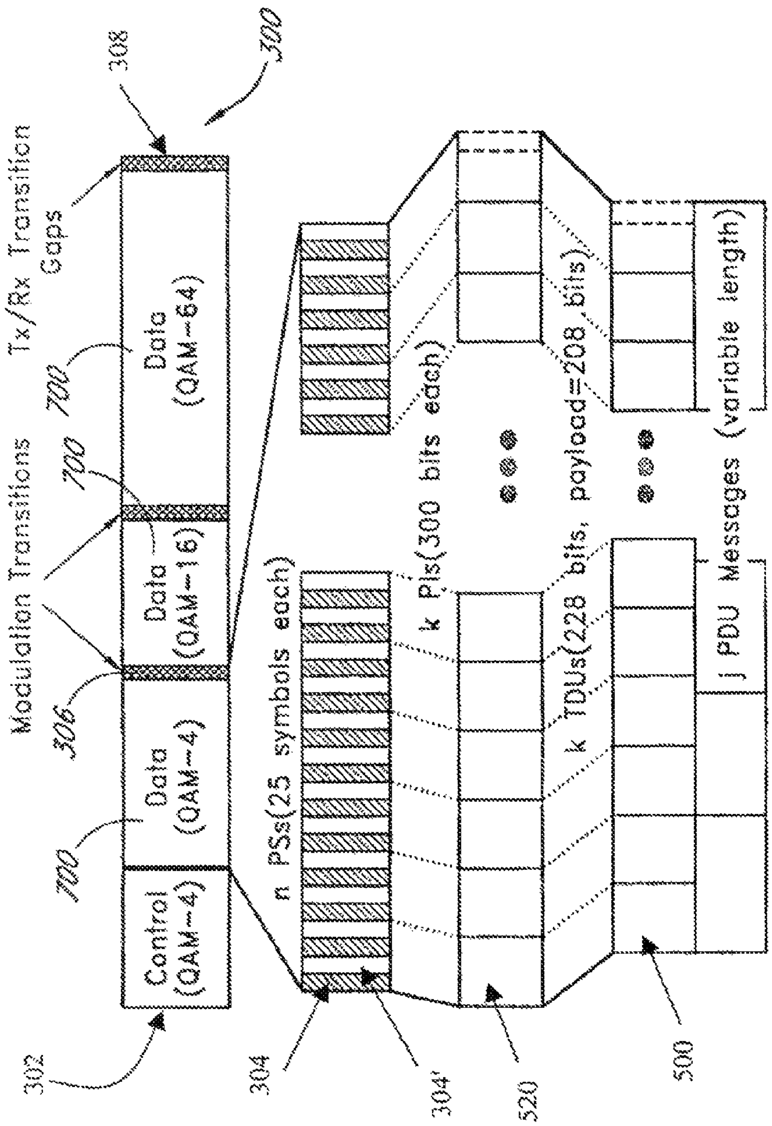

Tx/Rx Transition Gaps
Modulation Transitions
300
308
700
306
302
304
304'
520
500
Control (QAM-4)
Data (QAM-4)
Data (QAM-16)
Data (QAM-64)
n PSs(25 symbols each)
k PIs(300 bits each)
k TDUs(228 bits, payload=208 bits)
j PDU Messages (variable length)

Figure 5

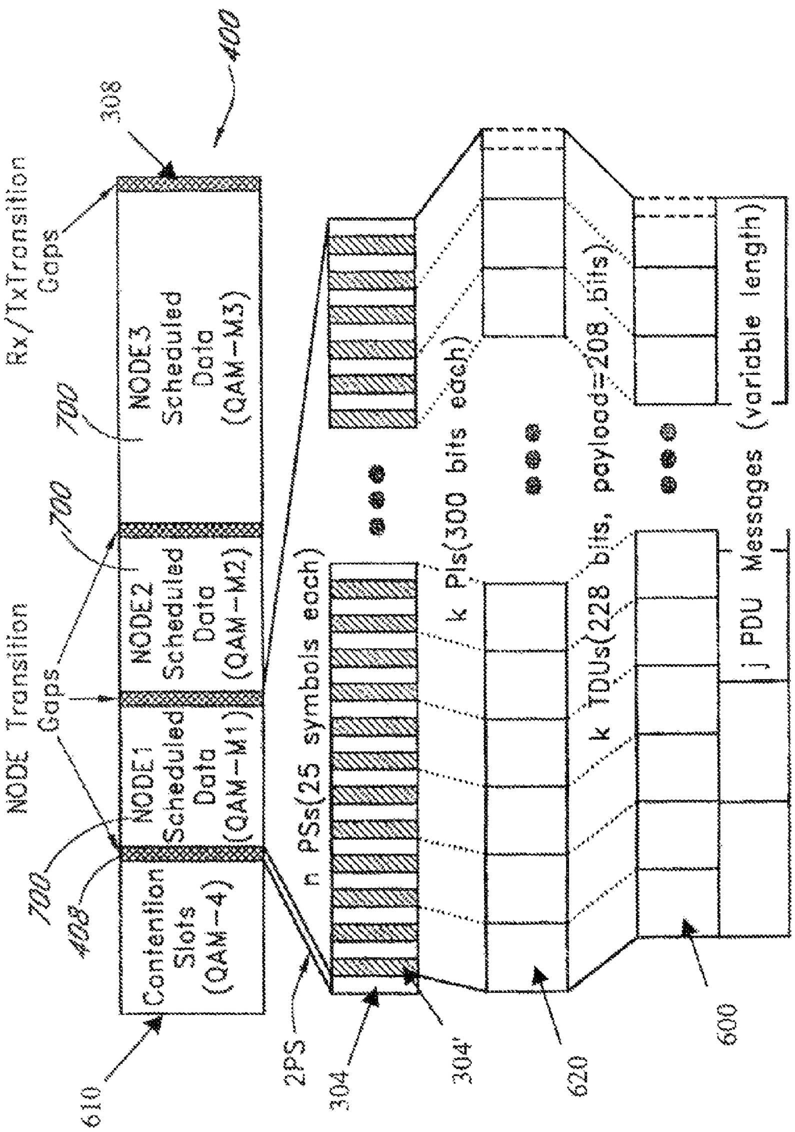
400
308
Rx/TxTransition Gaps
NODE Transition Gaps
700
408
610
Contention Slots (QAM-4)
NODE1 Scheduled Data (QAM-M1)
NODE2 Scheduled Data (QAM-M2)
NODE3 Scheduled Data (QAM-M3)
n PSs(25 symbols each)
k PIs(300 bits each)
k TDUs(228 bits, payload=208 bits)
j PDU Messages (variable length)
2PS
304
304'
620
600

Figure 6

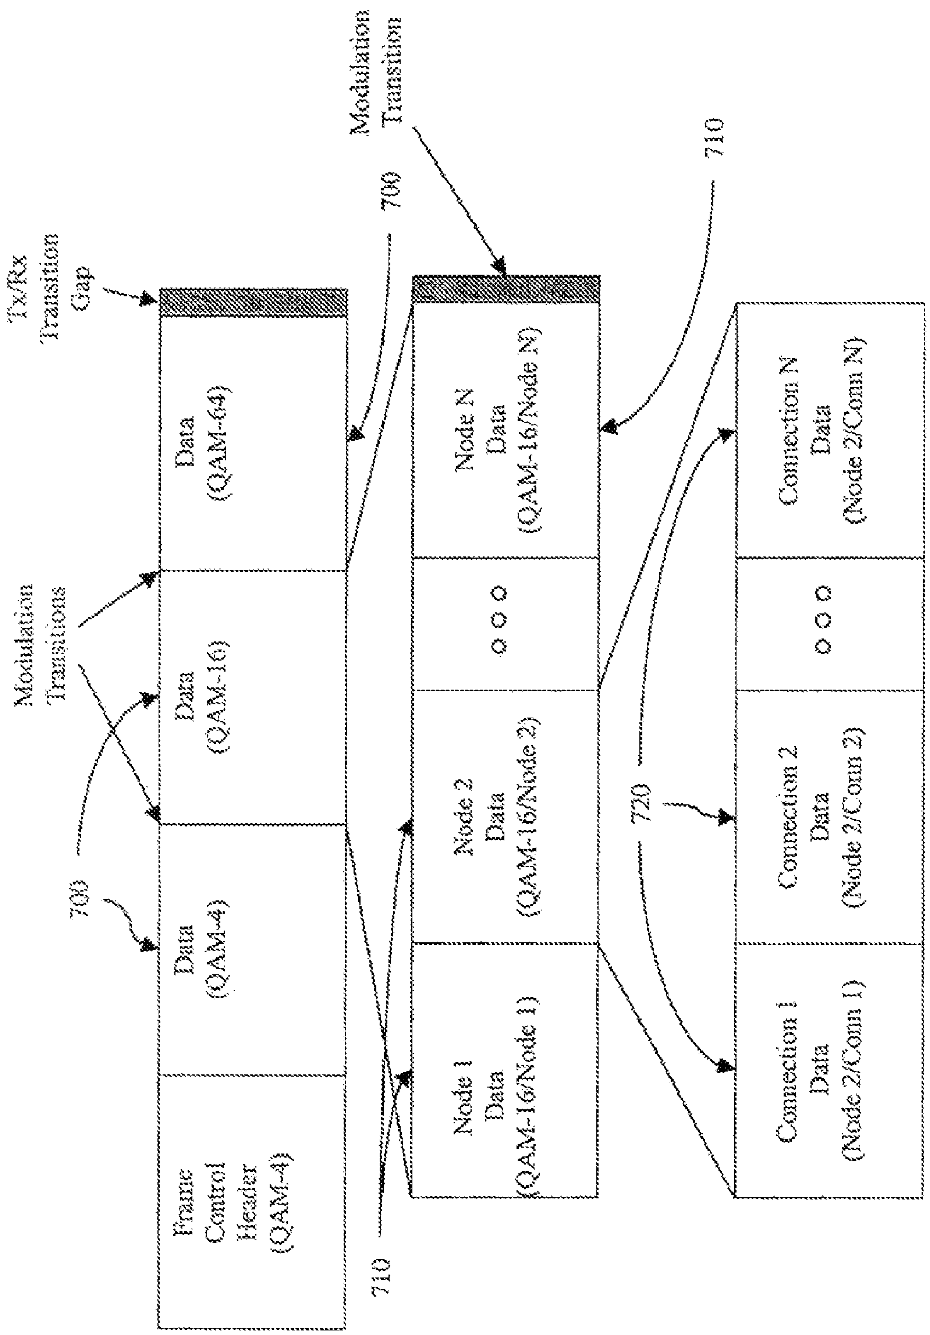
Tx/Rx Transition Gap
Modulation Transition
700
710
Data (QAM-64)
Node N Data (QAM-16/Node N)
Connection N Data (Node 2/Conn N)
Modulation Transitions
Data (QAM-16)
Node 2 Data (QAM-16/Node 2)
720
Connection 2 Data (Node 2/Conn 2)
700
Data (QAM-4)
Node 1 Data (QAM-16/Node 1)
Connection 1 Data (Node 2/Conn 1)
Frame Control Header (QAM-4)
710

Figure 7

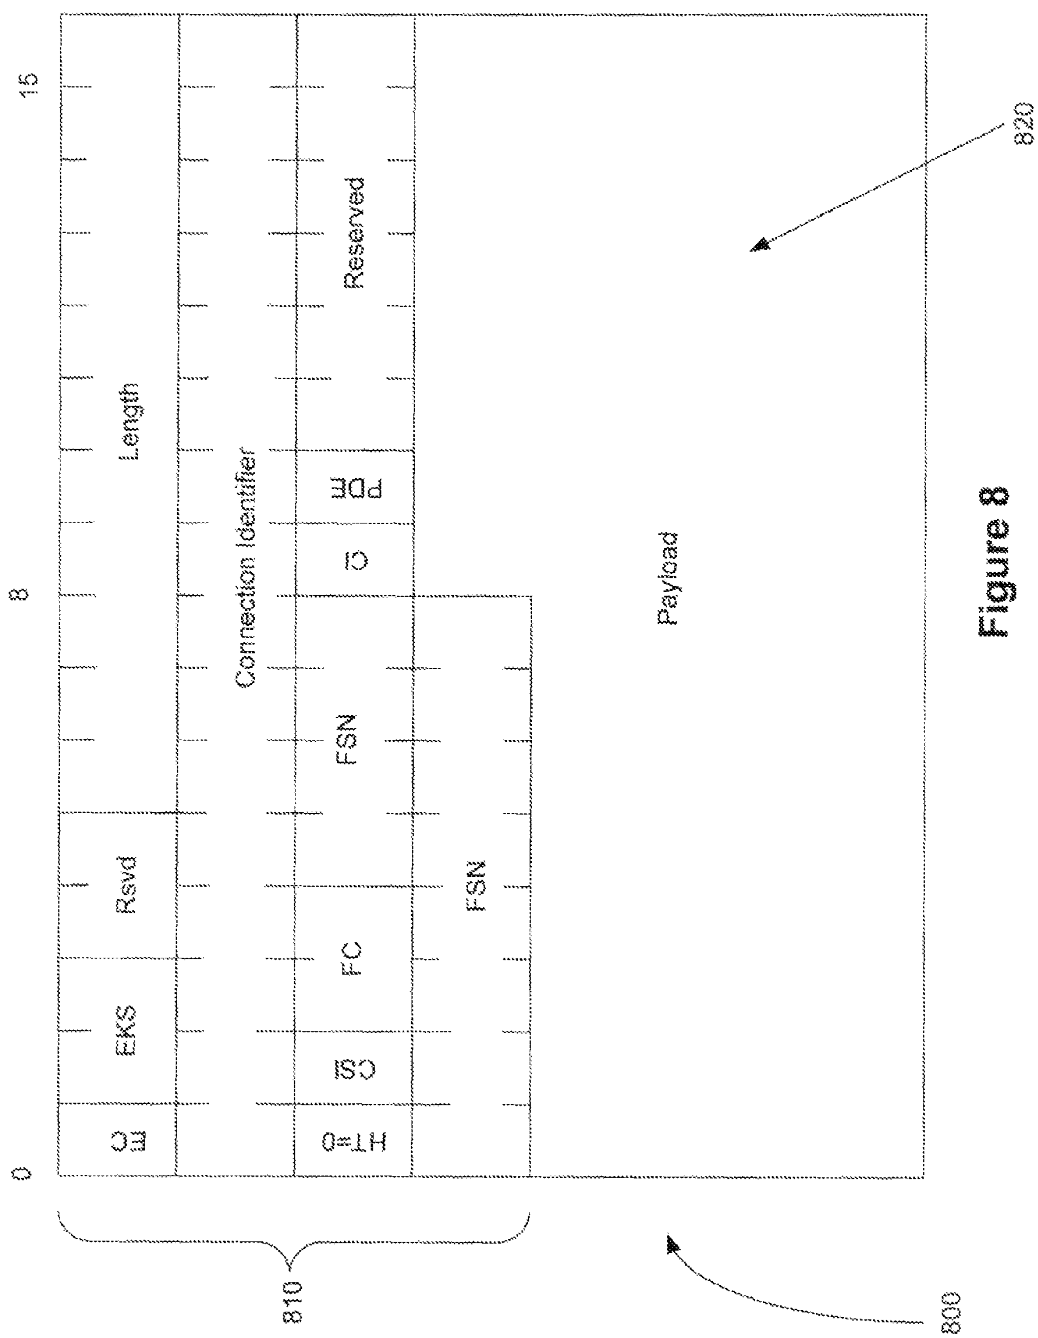
0
8
15
EC
EKS
Rsvd
Length
Connection Identifier
HT=0
CSI
FC
FSN
CI
PDE
Reserved
FSN
Payload
810
800
820

Figure 8

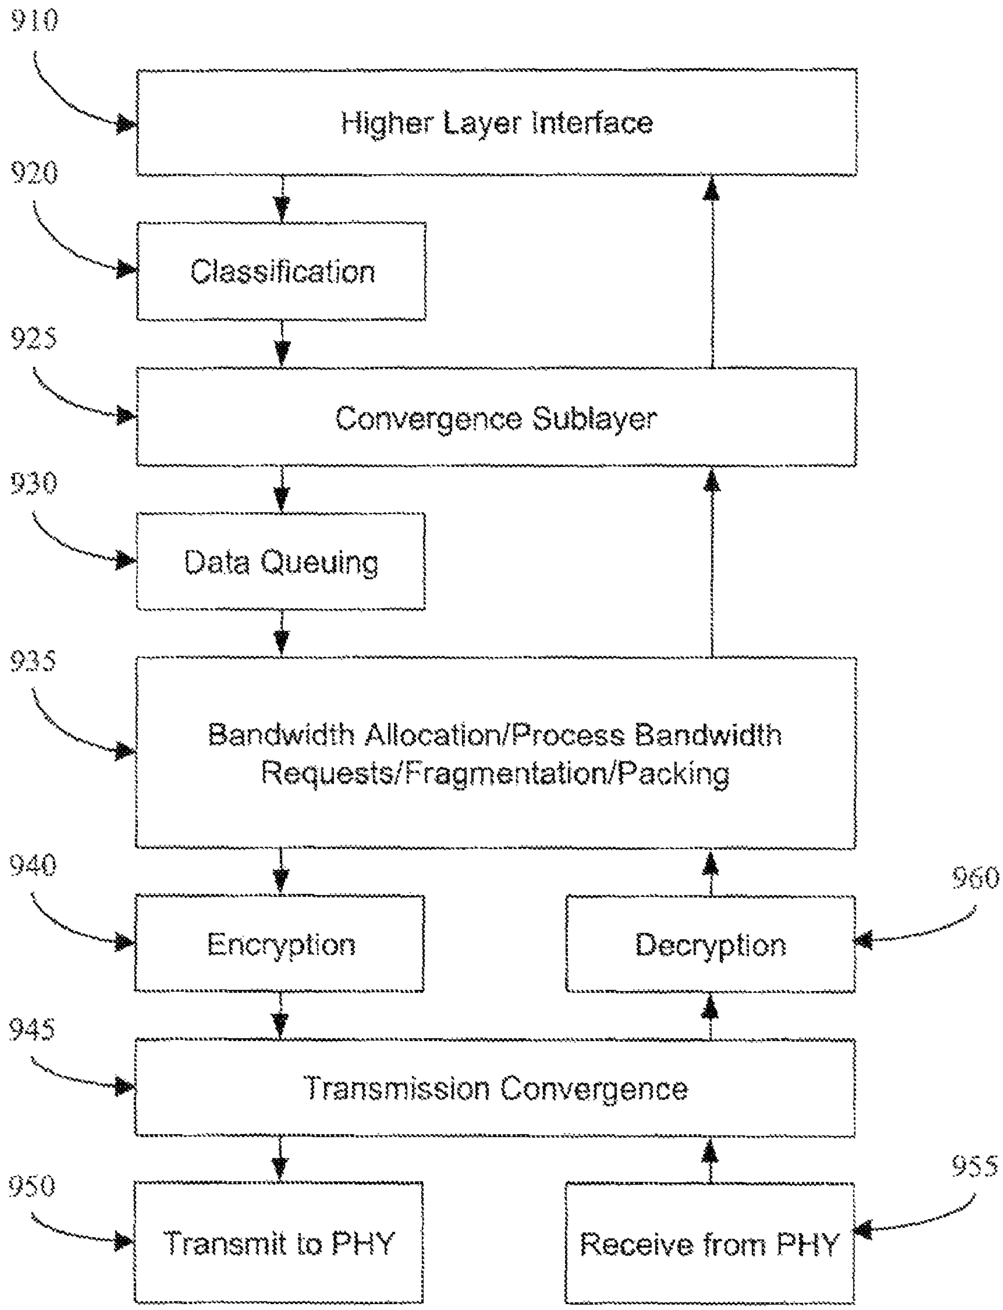
910
Higher Layer Interface
920
Classification
925
Convergence Sublayer
930
Data Queuing
935
Bandwidth Allocation/Process Bandwidth Requests/Fragmentation/Packing
940
Encryption
960
Decryption
945
Transmission Convergence
950
Transmit to PHY
955
Receive from PHY

Figure 9

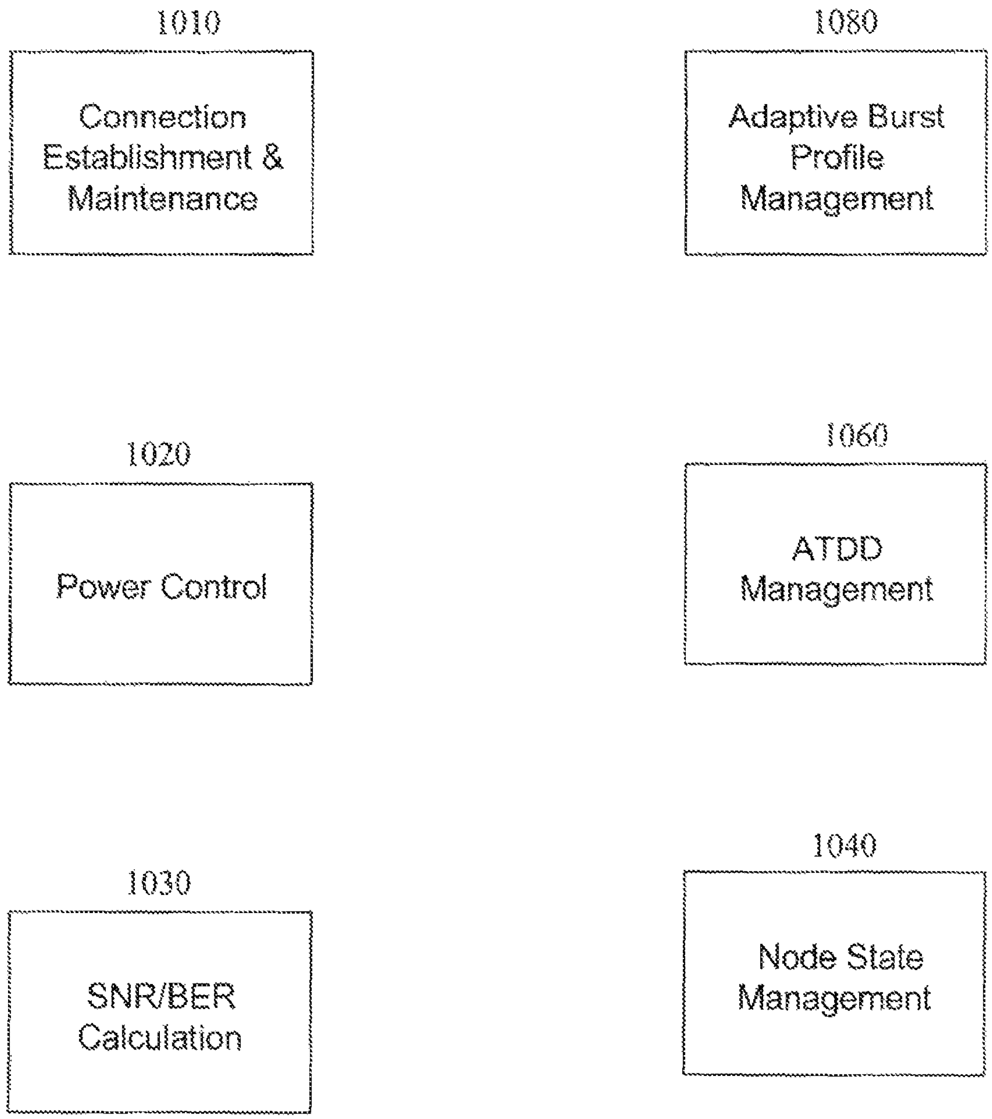
1010
Connection Establishment & Maintenance
1080
Adaptive Burst Profile Management
1020
Power Control
1060
ATDD Management
1030
SNR/BER Calculation
1040
Node State Management

Figure 10

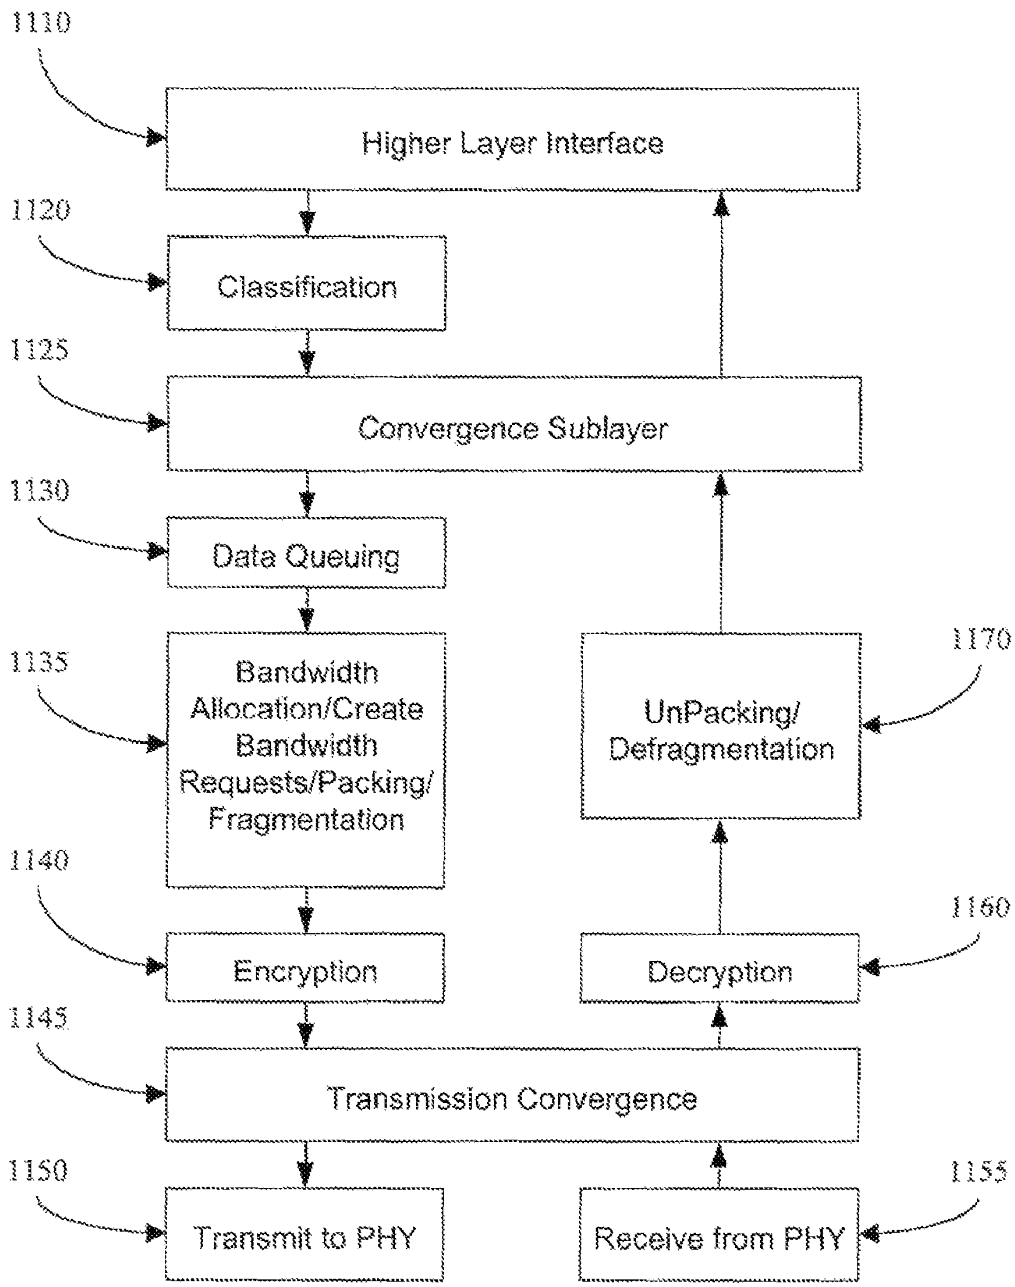

1110
Higher Layer Interface
1120
Classification
1125
Convergence Sublayer
1130
Data Queuing
1135
Bandwidth Allocation/Create Bandwidth Requests/Packing/ Fragmentation
1170
UnPacking/ Defragmentation
1140
Encryption
1160
Decryption
1145
Transmission Convergence
1150
Transmit to PHY
1155
Receive from PHY

Figure 11

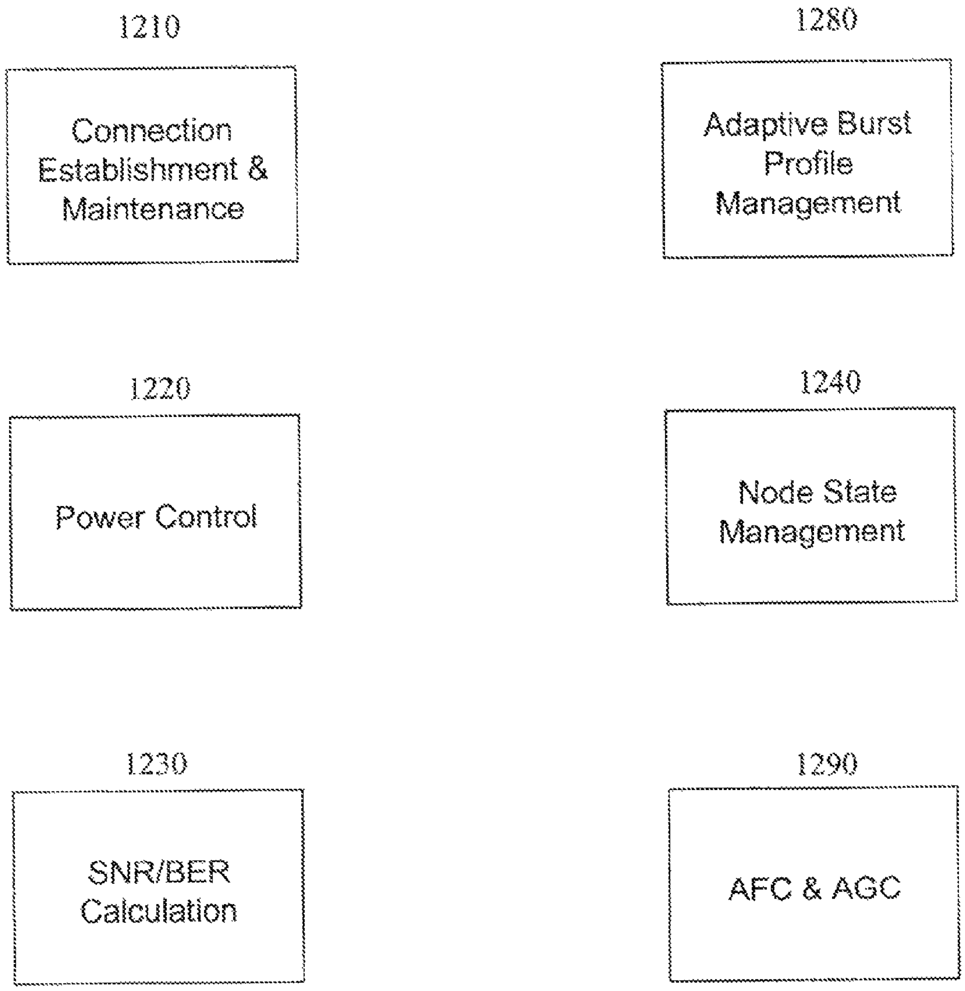
1210
Connection Establishment & Maintenance
1280
Adaptive Burst Profile Management
1220
Power Control
1240
Node State Management
1230
SNR/BER Calculation
1290
AFC & AGC

Figure 12

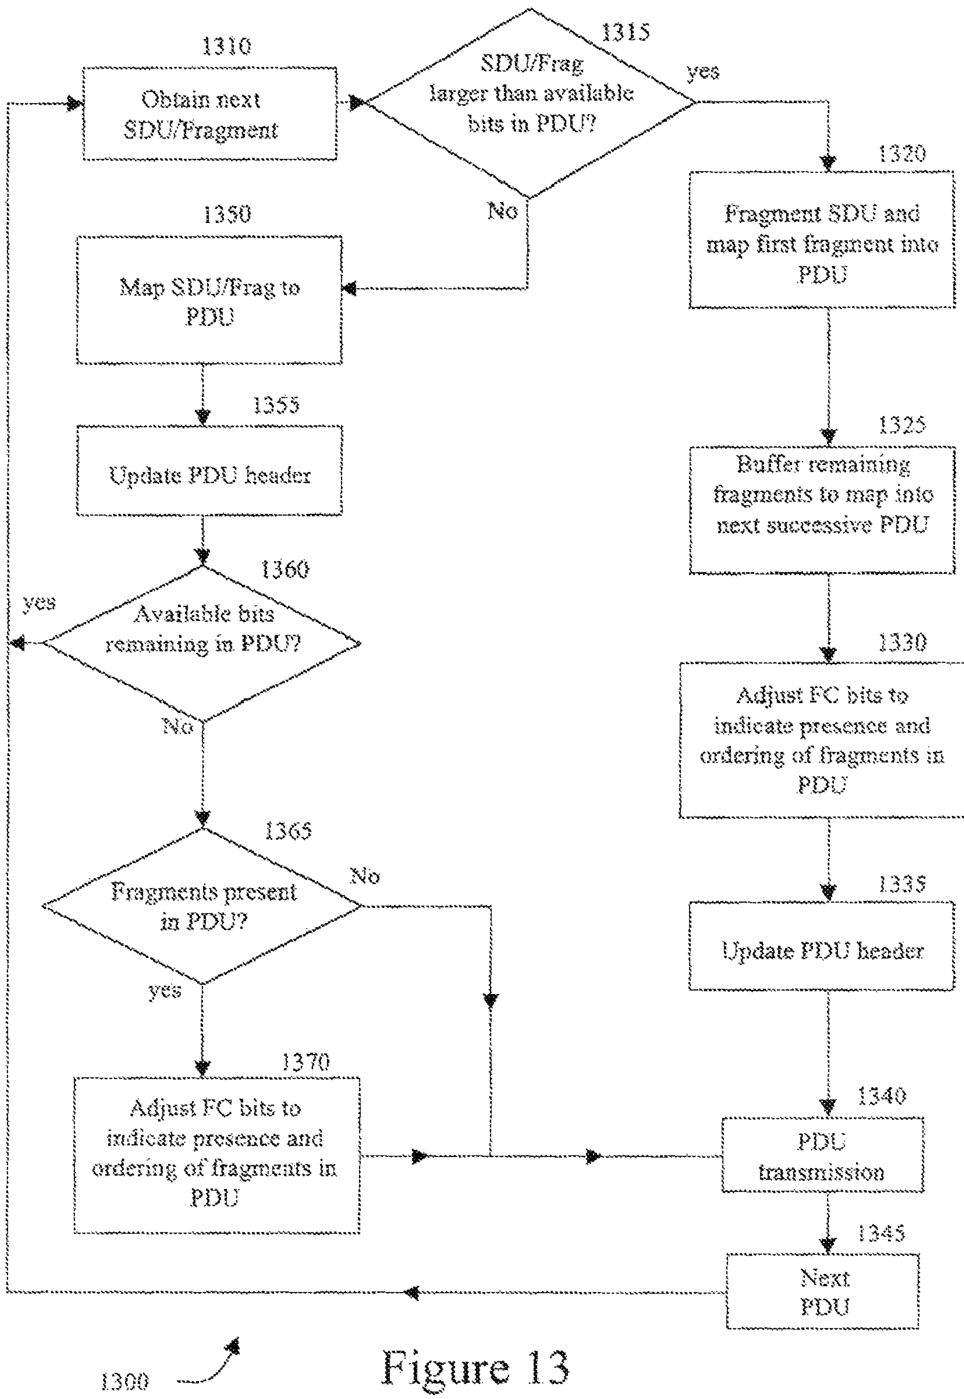
1310
Obtain next SDU/Fragment
1315
SDU/Frag larger than available bits in PDU?
yes
No
1320
Fragment SDU and map first fragment into PDU
1350
Map SDU/Frag to PDU
1355
Update PDU header
1325
Buffer remaining fragments to map into next successive PDU
1360
Available bits remaining in PDU?
yes
No
1330
Adjust FC bits to indicate presence and ordering of fragments in PDU
1365
Fragments present in PDU?
No
yes
1335
Update PDU header
1370
Adjust FC bits to indicate presence and ordering of fragments in PDU
1340
PDU transmission
1345
Next PDU
1300

Figure 13

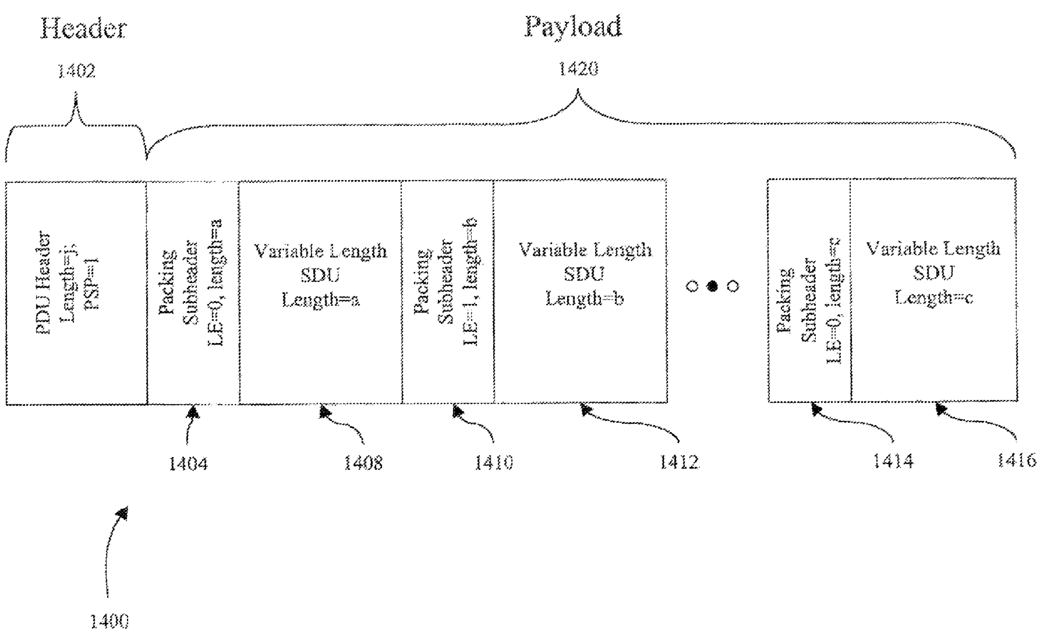

Header
1402
Payload
1420
PDU Header Length=j; PSP=1
Packing Subheader LE=0, length=a
Variable Length SDU Length=a
Packing Subheader LE=1, length=b
Variable Length SDU Length=b
Packing Subheader LE=0, length=c
Variable Length SDU Length=c
1404
1408
1410
1412
1414
1416
1400

Figure 14

PACKING SOURCE DATA PACKETS INTO TRANSPORTING PACKETS WITH FRAGMENTATION

RELATED APPLICATIONS

This application is a continuation application of U.S. application Ser. No. 10/053,179 filed Jan. 15, 2002; which claims priority to U.S. Provisional Application No. 60/262,005, filed on Jan. 16, 2001, all of which are incorporated herein by reference in their entirety.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to packet data communications systems, and reformatting data in such systems before transmitting the data through a link.

2. Description of the Related Art

Data communications systems typically transfer data from a source to an end user by routing the data in packets through communications links. All links have physical limits on their data-carrying capacity, or bandwidth. It is a constant pursuit to most efficiently utilize the finite capacity of any communications link in an effort to increase data throughput.

Many varieties of communications systems exist with a variety of different protocols governing their transmission of data. The information transmitted in many of these systems is transmitted in discrete packets of data. For each system these packets may be of a standard length or may vary in length as the needs of the users dictate, but the format of the packets are generally unique to the protocol utilized. Data packets utilizing a particular protocol and format may be referred to as service data units, or SDUs. An exemplary format, Internet Protocol or EP format, permits flexibility in the routing of data between a source and a destination, while other formats may convey voice data with limits on time delays, so as to ensure that the voice data can be reconstructed with adequate fidelity at the receiving end. It is often desirable for data in various formats to utilize the same data links as part of their transmission paths. This is particularly true for links directed to solving the problem of connecting end users to the various communications networks that are the source of data sought by those users, known as the "last mile" problem. Solutions for the "last mile" problem tend to attempt to satisfy, to the greatest extent possible, the needs of the users, while supporting the various protocols and packet formats that the data may utilize. While these solutions often involve various processes, it is often true that most communications links utilize a specific data packet format and protocol of their own to most efficiently utilize those links; the protocol data packets utilized by these links may be referred to as Protocol Data Units or PDUs. It is an ongoing effort in the data communications industry to maximize the efficiency of communications links having a finite bandwidth while maintaining the integrity of the protocol and format of the SDUs being transported by those links.

SUMMARY OF THE INVENTION

The systems and methods have several features, no single one of which is solely responsible for its desirable attributes. Without limiting the scope as expressed by the claims which follow, its more prominent features is now discussed briefly. After considering this discussion, and particularly after reading the section entitled "Detailed Description of the Preferred Embodiments" one will understand how the features of the system and methods provide several advantages over traditional communications systems.

In one aspect, the invention relates to a system and method of formatting data arriving in SDUs of various formats into different packets, having a PDU packet format, for transport across a communications link, comprising packing one or a plurality of fragments of arriving SDUs or whole SDUs into single PDU packets.

Within the above aspects, the plurality of SDUs may have different lengths, with a length of at least some of the SDUs reflected in respective packing subheaders. The packing subheaders may be made contiguous with the SDUs whose length within the PDU they reflect, or in another aspect may be separated from the SDUs. Other aspects include the foregoing systems or methods, further providing at least two fragmentation control bits in a header of the PDU. The fragmentation control bits may indicate absence or the presence and orientation of any fragments in the PDU.

Another aspect can be embodied in any packeted data communications network. For example, the aspect may be embodied in a broadband wireless link that connects a plurality of end users to various networks.

In another aspect, a millimeter wavelength wireless RF channel communications system connecting single base stations to a plurality of relatively proximate nodes can be utilized. A network of such base stations with their surrounding nodes can provide communications services over a large area, such as a city. Such a system is representative of a variety of communications links having a limited communications media which must be shared by a plurality of different entities. Such systems may include wire connected information distribution systems such as, for example, Dial-up or DSL systems, visible light spectrum data transmission systems, and microwave data transmission systems among others.

In yet another aspect, a method is disclosed of packing data prior to transferring it through a communications link while utilizing the advantages of fragmentation of SDU packets, and coordinating the two methods to optimize the advantages of each. Incoming SDU data packets that have been formatted according to a first, second or other protocol, such as ATM standard format or Internet Protocol (IP) or any other protocol, are compressed and reformatted and then conveyed over one or more links in accordance with a second PDU protocol, such as, for example, variable length packet MAC protocol. After it has traversed the link, the data may be reconstructed back into the first protocol format so that the data modifications performed by the link will be transparent to the receiving node or user.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. **1** is a high level illustration representing an overall communications network and system.

FIG. **2** is a high level functional block diagram of an exemplary base station.

FIG. **3** is a high level functional block diagram of an exemplary node.

FIG. **4** is an illustration of the breakdown of a frame in communications systems utilizing frames.

FIG. **5** is an illustration of the Downlink mapping of messages from PHY elements to PDUs in one embodiment.

FIG. **6** is an illustration of the Uplink mapping of messages from PDUs to PHY elements in one embodiment.

FIG. **7** is an illustration of an information hierarchy from sector transmission to modulation group information to connection breakdown in an exemplary transmission link frame.

FIG. **8** is an illustration representing a sample PDU header and illustrating the various fields the PDU header might have.

FIG. **9** is a functional block diagram of components that process and transfer data in an exemplary base station communications processor.

FIG. **10** is functional block diagram of components that handle control functions in an exemplary base station communications processor.

FIG. **11** is a functional block diagram of components that process and transfer data in an exemplary node communications processor.

FIG. **12** is functional block diagram of components that handle control functions in an exemplary node communications processor.

FIG. **13** is a flow chart illustrating the steps of an exemplary process for forming PDUs from incoming SDUs.

FIG. **14** is an illustration demonstrating the relationship between the PDU header, the payload and the PDU subheader.

DETAILED DESCRIPTION

Embodiments of the invention are now described with reference to the accompanying figures, wherein like numerals refer to like elements throughout. The terminology used in the description presented herein is not intended to be interpreted in any limited or restrictive manner simply because it is being utilized in conjunction with a detailed description of certain specific embodiments of the invention. Furthermore, embodiments of the invention may include several novel features, no single one of which is solely responsible for its desirable attributes or which is essential to practicing the inventions herein described.

Communications networks often need to transport data in a variety of different formats. It is therefore useful if communications systems that provide links within an overall communications network are able to accept data in a first or second format such as ATM or IP, or many other formats such as voice communications, Ti, El, or any other format common in the art. However, it is also useful for many communications system links to have a particular preferred format for all transported data in order to most efficiently utilize the link capabilities. To harmonize these two advantages, it is often desirable that a linking or transporting communications system accepts data in many formats and converts the data into a transporting format for transport across the link. At the far end of the link, the reformatted data may be returned to its original format. In this process of reformatting, incoming data arrives as SDUs (Service Data Units), which may be in any of the above mentioned formats, and will be converted to PDUs (Protocol Data Units) having a format desirable for a linking communications system. It is desirable to fit the SDUs efficiently into the PDUs to enhance the data carrying capability of the communications system link transporting the incoming data SDUs.

An example of a communications system that provides links for use within overall communications networks is described in copending U.S. patent application Ser. No. 09/702,293, entitled "COMPRESSION OF OVERHEAD IN LAYERED DATA COMMUNICATION LINKS," filed Oct. 30, 2000 (the '293 application), which is hereby incorporated by reference. The methods described herein may be employed with the system modules described in the '293 application to form an improved system for transporting data across a communications link. Appropriate functional modules described in the '293 application employing the specific methods of packing described below form a system and apparatus for packing SDUs into transport system PDUs that have the capacity to carry SDU fragments, which are portions of SDUs. Fragmenting is a method of partitioning a packet of data into two or more smaller pieces to be conveyed over a communications link that utilizes packets; and it is accomplished when the packet is too large for the existing bandwidth of a current frame or communications cycle.

FIG. **1** is a high level illustration representing an overall communications network and system. FIG. **1** depicts a representative network for transmitting data packets from a data source to end-users and vice versa. While FIG. **1** depicts a system utilizing a wireless link between a base station **12** and nodes **16**, this is only exemplary and any transmission link, such as electrical conductors, RF waves, microwaves, optical fiber conductors, and point to point/multipoint light transmission links, can be used. The communications domain of the base station **12** in FIG. **1** is directional and is broken up into four sectors **14** with each sector **14** capable of containing multiple nodes **16**. Although the base station of FIG. **1** provides for four sectors **14**, the base station **12** may be non-directional, having no sectors **14**, or may have more or less than four sectors **14**.

Within each sector **14** or transmission area of the base station **12** there may be any number of nodes **16**. The base station **12** may utilize one or more modulation and error correction schemes with which to transmit and receive signals at varying degrees of reliability and bandwidth. The wireless system **10** exemplified in FIG. **1** may utilize Time Division Duplexing (TDD), Frequency Division Duplexing (FDD) or any other duplexing or multiplexing, as well as any other type of communications link modulation or segmenting scheme. For simplicity of explanation, a TDD system will be made reference to hereinafter. As illustrated, some nodes **16** may utilize QAM-4 while others utilize QAM-16 and QAM-64; but the illustrated division between the nodes **16** is only exemplary and the modulation scheme utilized by any particular base station **12** may depend upon the connection establishing and monitoring routine and protocol of the particular system **10**.

The highlighted sector **14** contains five nodes **16** with each node **16** serving multiple connections for users. The users may be a service network such as a LAN, WAN, Intranet, Ring Network or other type of network; or they may be a single user such as a work station. The base station **12** is advantageously connected to various data sources such as the internet, other communications networks or any number of data bases, or any other data source. Information is received by the base station **12** from the data source, is prepared for and transmitted across a data link to a node **16**, and is then directed to the appropriate connection for transmission to the appropriate user.

Information is advantageously passed in the opposite direction as well, from user to data source.

Within the sectors **14**, the downlink transmissions from the base station **12** are typically multiplexed, and each node **16** within a particular sector **14** can receive the same transmission from the base station **12**. Each node **16** may await its particular information indicated by some control means and then process only the information contained therein; or alternatively each node **16** may receive all of the data within its modulation group and discard any data not pertinent to the users on its connections. Nonetheless, each node **16** has a distinct "virtual" connection, or communications link, within its sector **14**. A link conveys that part of the downlink transmissions within the sector **14** from the base station **12** that contains data for the particular node **16**, and also conveys uplink transmissions from the particular node **16** to the base

station **12**. Nodes **16** in other sectors **14** likewise communicate to base station **12** through links that are virtual connections within the transmissions of their particular sector **14**. The transmissions of different sectors **14** are independent of each other. "Sectorized" transmission permits spectrum reuse within a narrow area, thus providing more bandwidth to service particular users.

FIG. **2** is a high level functional block diagram of an exemplary base station and illustrates the functional modules that may be used. The term "module," as used herein, means, but is not limited to, a software or hardware component, such as an FPGA or ASIC, which performs certain tasks. A module may advantageously be configured to reside on an addressable storage medium and configured to execute on one or more processors. Thus, a module may include, by way of example, components, such as software components, object-oriented software components, class components and task components, processes, functions, attributes, procedures, subroutines, segments of program code, drivers, firmware, microcode, circuitry, data, databases, data structures, tables, arrays, and variables. The functionality provided for in the components and modules may be combined into fewer components and modules or further separated into additional components and modules. Additionally, the components and modules may advantageously be implemented to execute on one or more computers.

A base station **12** may comprise a communications processor **20**, a modem **22**, an antenna **24**, an input/output (I/O) control **26**, and an IF/RF converter **30**. These modules indicate certain functions but are not intended to indicate any particular architecture. In fact, the functions represented by these modules may be combined into a single module, into multiple modules or any combination thereof, with the illustration in FIG. **2** providing only an exemplary arrangement of one way in which to carry out those functions.

The communications processor **20**, or converter, fulfills many functions as described below and provides most of the control functions occurring within the base station **12**. During downlink transmissions, the data sources provide data in the form of SDUs to the base station **12** via the backhaul interface **28**, which forms the connection between the data sources and the base station **12**. The I/O control **26** controls the transfer of SDUs between the base station **12** and the data sources via the backhaul interface **28**. The I/O controller **26** transfers the SDUs from the backhaul interface **28** to the communications processor **20**, which among many other things, converts them to PDUs of a protocol format that is appropriate for the transmission link. The communications processor **20** transfers the PDUs to a modem **22**, which converts them to an intermediate transmission modulation or frequency for an RF link system, and passes them on to an IF-RF converter **30**. The IF-RF converter **30** converts intermediate frequency signals provided by the modem **22** to an appropriate frequency required for transmission before passing them to the antenna **24**. In a system that does not utilize an RF link, this step may not be necessary, or may take another form that is appropriate for that medium. The antenna **24** preferably receives signals from the IF-RF converter **30** at a radio frequency, or transmission frequency, and transmits them. For systems that do not utilize an RF link, other suitable transmission mechanisms are utilized. In other words, the function provided by the IF-RF converter **30** and antenna **24** may generally be thought of as that of a transmitter in any system, wireless or not.

During uplink transmissions, the antenna **24** receives RF signals from one or more nodes **16** and transfers them to the IF-RF converter **30**. The IF-RF converter **30** converts signals provided by the antenna **24** to an appropriate frequency range for the modem **22** to process. The general function performed by the antenna **24** and the IF-RF converter **30** may be thought of as that of a receiver in this or other systems. For nonRF systems, comparable modules would perform these functions to prepare the received signals for the modem **22**. The modem **22** demodulates the signal from the IF-RF converter **30** and transfers a digital signal comprised of PDUs to the communications processor **20**. The communications processor **20** receives the digital signal from the modem **22** and, among other things, converts the signal into the SDUs that the users had transmitted to the node **16**. The SDUs are then sent to the PO control **26** for transfer out of the base station **12**. The I/O control **26** transmits the SDUs to the appropriate data source via the backhaul interface **28**.

FIG. **3** is a high-level block diagram of the functional modules of an exemplary node. A node **16** may include a communications processor **32**, a modem **40**, an IF-RF converter **42**, an antenna **44**, and a connection interface **34** coupled to a plurality of user connections **36**. These modules indicate certain functions, but are not intended to indicate any particular architecture. The functions may be fulfilled by any particular module alone or in any combination. Alternatively, a single module may accomplish all of the functions.

During downlink transmissions, the PDUs are transmitted from the base station **12** to the node **16** and are received at the node **16** by the antenna **44**, for RF systems. Systems not using RF communications links would use a receiver having analogous receiving components. The antenna **44** converts the RF signals received into electronic signals which are then transferred to the IF-RF converter **42**. The IF-RF converter **42** converts the signals from the transmission frequency to an intermediate frequency and transfers those signals to the modem **40**. The modem **40** further demodulates the intermediate frequency signals into a digital signal that includes PDUs. The digital signal including PDUs that are then transferred to the communications processor **32**, or converter, which then converts the PDUs back into the SDUs that were sent to the base station **12** by the data source(s). The SDUs are then directed to the connection interface **34**, which directs the SDUs to the appropriate user connection **36**. The SDUs can then be directed to the appropriate users via the user connections **36**.

During uplink transmissions, information packages in the form of SDUs are provided by the user connections **36** to the connection interface **34**. The connection interface **34** is utilized by the communications processor **32** to control the transmission of SDUs to the node **16** and transfers the SDUs to the communications processor **32**. The communications processor **32**, among other things, converts the SDUs into the appropriate PDU format for the transmission link. The PDUs are then transferred from the communications processor **32** to the modem **40**, which modulates them onto an IF carrier signal. The modem **40** transfers the IF signal to an IF-RF converter **42**, which further converts the signal to the RF range that is appropriate for the antenna **44** or other transmitting mechanism. Again, if an RF communications link is not being utilized, the IF-RF converter **42** and the antenna **44** may be substituted by an appropriate transmitter function module. This function is that of a transmitter and any suitable transmitter may be used. The RF signal is then transmitted via the antenna **44** across the comminations link to the base station **12** for processing and transference to an appropriate data source as discussed above.

FIG. **4** is an illustration of the breakdown of a frame in communications systems utilizing frames. FIG. **4** shows a TDD frame and multi-frame structure **200** that may be used

by the communications system **10** of FIG. **1**. As shown in FIG. **4**, the TDD frame **200** is subdivided into a plurality of physical slots (PS) **204**, **204'**. In one embodiment, the TDD frame **200** is one millisecond in duration and includes 800 physical slots. Alternatively, the present invention can be used with frames having longer or shorter duration and with more or less PSs. Some form of digital encoding, such as the well-known Reed-Solomon (RS) encoding, convolutional encoding, or turbo code encoding, may be performed on the digital information over a pre-defined number of bit units referred to as physical layer information elements (PI). The modulation and/or the FEC type may vary within the frame and determines the number of PSs (and therefore the amount of time) required to transmit a selected PI. In one embodiment, data is referred to as being sent and received using three different modulation types, namely, QAM-4, QAM-16, and QAM-64.

In alternative embodiments, other modulation types, FEC types, or variation of a modulation or FEC type may be used. For example, an RS encoding system may use different variations of block sizes or code shortening; a convolutional encoding system may vary the code rate; and a turbo code system may use any block size, code rate, or code shortening. To aid periodic functions, multiple frames may be grouped into multiframes **206**, and multiple multi-frames **206** may be grouped into hyper-frames **208**. In one embodiment, an Adaptive Time Division Duplex (ATDD) system may be implemented. In ATDD mode, the percentage of the TDD frame allocated to downlink versus uplink is a system parameter that may change with time. In other words, an ATDD system may vary the ratio of downlink data to uplink data in sequential time frames.

FIG. **5** is an illustration of the downlink mapping of messages from PHY elements to PDUs in one embodiment. FIG. **5** shows one example of a TDD downlink subframe **300** that can be used by the base station **12** to transmit information to the plurality of nodes **16**. The downlink mapping illustrated in FIG. **5** can be performed by the communications processor in the base station, and is performed to map PDUs of varying lengths to the PSs utilized by a wireless communications system as described above with reference to FIG. **4**. As mentioned previously, in a TDD system, each time frame is divided into a downlink subframe and an uplink subframe. More specifically, during each frame (or other predetermined period), the downlink subframe is first transmitted from the base station **12** to all nodes **16** in the sector **14**, after which the uplink subframe is received by the base station **12** from particular nodes **16**. The downlink subframe **300** may be dynamic, such that it may be different in sequential time frames depending on, among other things, an uplink/downlink split determined by the communications processor **20**. In an FDD system, the time frame is not divided between uplink and downlink data. Instead, an FDD downlink subframe is an entire frame of downlink data on a first channel, and an uplink subframe is an entire frame of uplink data on a second channel. In a typical FDD system, the downlink subframe and uplink subframe may be transmitted simultaneously during the same predetermined period. Thus, in an FDD system both the base station **12** and the nodes **16** may receive and transmit at the same time, using different channels. In another embodiment, the downlink subframe and uplink subframe may not be transmitted at the same time, but still use different channels.

The downlink subframe **300** of FIG. **5** preferably comprises a frame control header **302** and a plurality of downlink data PSs **204**. The plurality of data PSs may be grouped by any combination of modulation type, FEC type, node index, and connection ID and may also be separated by associated modulation transitions (MTs) **306**. MTs separate differently modulated data, and a transmit/receive (Tx/Rx) transition gap **308**. MTs may be a gap, a period of time to allow for the transition from one modulation group to the next. Alternatively, that transition can occur at the boundary between the last PS of one modulation group and the first PS of the next modulation group. In any downlink subframe, any one or more of the differently modulated data blocks may be absent. In one embodiment, MTs **306** are 0 (“zero”) PSs in duration. The frame control header **302** may contain a preamble that can be used by the physical protocol layer (PHY) for synchronization and equalization purposes. The frame control header **302** also includes control sections for both the PHY and the PDU protocol controls. An FDD downlink subframe may be substantially identical to the structure of FIG. **5**, but without a Tx/Rx transition gap **308**.

The downlink data PSs **304**, **304'** are advantageously used for transmitting data and control messages to the nodes **16**. This data is preferably encoded (using a ReedSolomon encoding scheme, or other scheme for example) and transmitted at the current operating modulation used by the selected node **16**. In one embodiment, data is transmitted in a pre-defined modulation sequence, such as QAM-4, followed by QAM16, followed by QAM-64. The MTs **306**, if present, are used to separate the modulation schemes to synchronize the base station **12** and the nodes **16**. The PHY control portion of the frame control header **302** preferably contains a broadcast message to all of the nodes indicating the identity of the PS **304** at which the modulation scheme changes. The ordering of modulation groups in the transmission subframe illustrated in FIG. **5** is only an example and any ordering of modulation groups may be used; alternatively, the order may change from frame to frame depending on the needs of the system. Finally, as shown in FIG. **5**, the Tx/Rx transition gap **308** separates the downlink subframe from the uplink subframe. While the present embodiment illustrates the use of a gap to transition from uplink to downlink, systems may be equipped so as to identify the transition without the use of gaps.

FIG. **5** also shows an embodiment of a three-stage mapping process from a stream of variable length PDUs or user messages, to 228-bit TC Data Units (TDUs) **500**, otherwise known as a TC/PHY packets **500**, to 300-bit PIs and finally to 25-symbol PSs (PIs and PSs are described above with reference to FIG. **4**). The illustration and description of the process for conversion of SDUs into PDUs is discussed in detail later.

In one embodiment, a minimum physical unit that the system allocates is the 25-symbol PS **304**, **304'**. The minimum logical unit the exemplary system of FIG. **5** allocates may be a 208-bit (26-byte) payload of the 228-bit TC Data Unit (TDU) **500**. Other embodiments can be used that have different minimum quantities of the physical and logical units without departing from the scope of the present invention.

Alternatively, information mapping processes may take different steps between PDU formation and transmission. For instance, the last TDU (and PI) of a particular modulation may be shortened if there is not enough data to fill the entire TDU. This variability of the length of the last TDU and PI are illustrated by the dashed lines in the last TDU and PI in FIG. **5**; and the length of the last TDU and PI may be any length shorter than, or including, their respective full ordinary lengths.

The 228-bit TDU **500** may be encoded using the well-known Reed-Solomon coding technique to create the 300-bit PIs **520**. Bandwidth needs that do not require encoding, such as the various transition gaps, are preferably allocated in units of 1 PS. Bandwidth needs that require encoding (using a Reed-Solomon encoding scheme, for example) may be allo-

cated in TDUs **500**. Also, data for each modulation scheme, on the downlink, and each node's transmission, on the uplink, are advantageously packed and fragmented to form an integer multiple of TDUs **500** to create an integer multiple of PIs **520** or, alternatively, may be packed and fragmented into an additional fractional and shortened TDU to create a fractional and shortened PI. The number of PSs **304**, **304'** required to transmit a PI **520** may vary with the modulation scheme used. An exemplary system for mapping PDUs to the PHY, and vice versa, is described in detail in commonly assigned Patent Cooperation Treaty Application Number PCTUS00/29687, entitled METHOD AND APPARATUS FOR DATA TRANSPORTATION AND SYNCHRONIZATION BETWEEN MAC AND PHYSICAL LAYERS IN A WIRELESS COMMUNICATION SYSTEM (the '687 application"), which is hereby incorporated by reference. The mapping from PDU to PHY in the '687 application discloses a means of converting PDUs to a form appropriate for transmission by a wireless link. Similar systems can be used for embodiments utilizing different communications links.

FIG. **6** is an illustration of the uplink mapping of messages from PDUs to PHY elements in one embodiment. The uplink of data from upper layers to the PHY layer may occur in the communications processors of the various nodes served by each base station. FIG. **6** shows an example of an uplink subframe **400** that may be adapted for use with the data transportation and synchronization process. The nodes **16** use the uplink subframe **400** to transmit information (including, for example, bandwidth requests) to their associated base stations. There may be three or more main classes of control messages that are transmitted by the nodes during the uplink subframe **400**. Examples include: (1) those that are transmitted in contention slots reserved for node registration (Registration Contention Slots); (2) those that are transmitted in contention slots reserved for responses to multicast and broadcast polls for bandwidth allocation (Bandwidth Request Contention Slots); and (3) those that are transmitted in bandwidth specifically allocated to individual nodes (node Scheduled Data Slots).

During its scheduled transmission time, a node typically transmits in a selected modulation that can be selected based upon, for example, the effects of environmental factors on transmission between that node and its associated base station. The uplink subframe **400** includes a plurality of node transition gaps (NTGs) **408** that serve a function similar to that of the MTs described above. That is, the NTGs **408** allow for the transition from one node to the next during the uplink subframe **400**. In one embodiment, the NTGs **408** are 2 physical slots in duration. A transmitting node **16** may transmit a 1 PS preamble during the second PS of the NTG **408** thereby allowing the base station to synchronize to the transmission of the new node. In other embodiments, node transitions may alternatively occur at the transition between the last PS of one node's uplink transmission and the first PS of the next node's uplink transmission. One embodiment utilizes a system similar to that described in the '687 application for transmitting data from nodes to a base station; and this system and method should be understood to advantageously utilize such a system.

As illustrated in FIG. **6**, an uplink subframe **400** may comprise uplink contention access slots **610** as well as data slots **700**. The uplink contention access slots **610** may include registration contention slots (not shown) and bandwidth request contention slots (not shown). The uplink subframe **400** may begin with optional registration contention slots, or alternatively, the registration contention slots may be located at other points of the uplink subframe such as in the middle or

at the end. Some registration contention slots are allocated periodically to the physical slots for use during node registration. In one embodiment, registration messages are preceded by a 1 PS preamble and are sent alone. Also, other PDU control messages are not packed into the same PDU. The bandwidth request contention slots may be allocated for responses to multicast and broadcast polls for bandwidth requirements. In one embodiment, the bandwidth request messages, when transmitted in the bandwidth request contention period, may be proceeded by a 1 PS preamble. Nodes may pack additional bandwidth requests for other connections into the same PDU.

FIG. **6** also shows the mapping of the scheduled portion of the uplink subframe **400**. Within the subframe **400**, the TC/PHY packets **700** can be grouped by nodes. All transmissions from an individual node **16**, other than bandwidth requests transmitted in bandwidth request contention slots, may be transmitted using the same modulation scheme. In one embodiment, each node's transmission is packed and fragmented to be an integer multiple of a TDUs **600** to provide an integer multiple of PIs **620** after coding. In an alternative embodiment, if the bandwidth requested for pending uplink data does not require the entire last TDU, the bandwidth may be allocated such that the last TDU is shortened, resulting in a shortened PI. Again, this variability of the length of the last TDU and PI are illustrated by the dashed lines in the last TDU and PI in FIG. **6**; and the length of the last TDU and PI may be any length shorter than, or including, their respective full ordinary lengths. The uplink and downlink mapping provides a mechanism for the higher communications protocol layers to transport data to the PHY layer.

By using such a data transportation and synchronization technique, scheduled uplink and downlink data is transported and synchronized between the PDU processing layer (discussed below as item **935** in FIG. **9**) and the physical layer. The scheduled uplink and downlink data are preferably transported within the uplink subframe **400** and the downlink subframe **300**, respectively, based upon the modulation scheme used by the nodes **16**. Uplink mapping of PDUs to PHY elements may be performed according to the three stage process of PDU to TDU **600**, then from TDU **600** to PI **620**, then from PI **620** to PS. However, it is to be understood that there are numerous processes that are analogous and similar that may have more or less steps and may be used. Again, the process of converting SDUs to PDUs is described in detail later and the mapping described here provides understanding of how PDUs may be allocated to PIs in one wireless embodiment.

FIG. **7** is an illustration of an information hierarchy from sector transmission to modulation group information to connection breakdown in an exemplary transmission link frame. FIG. **7** illustrates the way information might be organized in messages sent from the base station to the nodes. As mentioned above, the base station transmits messages to the nodes containing three main categories of information; 1) a frame control header **302**, containing information to the nodes concerning the handling of the data, 2) the data **700** being conveyed from the data sources to the end users, and 3) the gaps **306**, **308** that separate the different sections of the transmission. The data can be broken down into the different modulation groups such as the system illustrated by FIG. **7** wherein those are the QAM-4, QAM-16, and QAM-64 modulation groups. A system may have more or less modulation groups. Again, the ordering of the modulation groups illustrated in FIG. **7** is only exemplary and any ordering or a shifting order may be utilized. FIG. **7** also illustrates a Tx/Rx transition gap, or a period of time for transition from the downlink subframe

to the uplink subframe. It should be noted however that the transition from downlink to uplink may alternatively occur at the boundary between the last PS of the downlink subframe and the first PS of the uplink subframe.

For each modulation group, the data **700** contains information for each node, the node data **710**. As mentioned above, those nodes **16** may either download all the information or just their assigned information. The illustration of FIG. **7**, depicting the information for each node ordered in a similar manner as the modulation group, is merely exemplary as well. The information for a particular node may also be spread throughout a modulation group downlink, or it may only be discretely located in one portion, or it may be discretely located in several portions. The information intended for each node **16** contains information to be distributed to the end users or services served by the connections of that node; this is identified in FIG. **7** as connection data **720**. The connection data **720** includes the information to be transmitted to the users or services as well as control information the node uses to identify to which of its connections each packet of information should be directed. Thus, the node can ensure that each of the packets of information that it receives is directed to the appropriate connection to reach the intended end user or service. In this manner, information transmitted by the base station **12** may logically be divided into modulation data groups **700**, and further into node data **710**, and further yet into connection data **720**. It should be noted that all of the modulation data groups **700**, node data **710** groups and connection data groups can be variable in size and may vary from frame to frame as well.

FIG. **8** is an illustration representing an exemplary PDU and illustrating the various fields the PDU header might have. FIG. **8** shows the format of one downlink PDU **800**. Although specific fields, field lengths, and field configurations are described with reference to FIG. **8**, those skilled in the communications art shall recognize that alternative configurations may be used including additional or fewer fields. In several embodiments, the communications processors of both the base station and the nodes create PDU payloads and PDU headers to be transmitted and retrieve SDUs from received PDUs. An exemplary downlink PDU format **800** may include a standard downlink PDU header **810** and a variable length PDU payload **820**. The downlink PDU header **810** of one embodiment comprises 13 different fields that measure 7 bytes in total length. The downlink PDU header **810** illustrated in FIG. **8** begins with an encryption control (EC) field. In certain embodiments, the EC field is set to a logical zero if the payload is encrypted; and it is set to a logical one if the payload is not encrypted. The EC field is followed by an encryption key sequence (EKS) field that provides information about the encryption used, if encryption is utilized. A reserved field (Rsvd) may follow the EKS field. The Rsvd field provides for future expansion of the PDU header fields. The Rsvd field is followed by a length field (Length). The Length field indicates the length of the PDU header and any data may contained in the PDU payload. The Connection Identifier field follows the Length field and provides identification information to the base station and the nodes. The Connection Identifier field identifies the destination to which each PDU is to be delivered.

A header type field (HT) follows the Connection Identifier field and indicates whether the header is a standard header or a bandwidth request header. The HT field is followed by a convergence sub-layer identification field (CSI) that provides information so that the communications processor may determine for which sub-layer, among equivalent convergence sub-layer peers, the PDU is intended. The CSI field is followed by a fragmentation control field (FC) and a fragmentation sequence number field (FSN). These two fields allow the communications processor to fragment SDUs to most efficiently utilize the payload of the PDU. The FC and FSN fields indicate the presence and orientation in the payload of any fragments of SDUs. The type of fragments present in the payload and the orientation of those fragments in the payload may vary. For example, fragmentation may result in a fragment that is the first fragment of an SDU, a continuing fragment of an SDU or the last fragment of an SDU.

In one embodiment, a particular SDU may be large enough to require several PDUs to transport it across the communications link. This SDU may require a first PDU to convey the first fragment of the SDU, several other PDUs to convey continuing fragments, and a final PDU to convey the last fragment of the SDU. In this embodiment the FC and FSN bits would indicate that the last part of the payload of the first PDU is a fragment. The FC and FSN bits of the middle PDUs would indicate that they contain continuing fragments in their payload and the last PDU would have FC and FSN bits to indicate that it contains the last fragment of a continuing SDU. The FC and FSN bits would also indicate where each of the fragments is in their respective PDU payload. For instance, the first fragment may be located at the end of the first PDU's payload, while the continuing fragments may take up the whole payload of their associated PDU and the last fragment may be at the beginning of the last PDU payload. It should be noted however, that more FC or FSN bits may be utilized to indicate any combination of types of fragments present in a PDU and their locations with respect to any whole SDUs in the payload and that the FC and FSN bits do not have to be located in the header but may be located elsewhere.

A CRC indicator field (CI) follows the FC and FSN fields to indicate whether or not CRC is appended to the payload. A packet discard eligibility field (PDE) can also be used and may provide information regarding the payload in a situation where there is congestion. In a congestion situation the wireless communications system may first discard packets indicating discard eligibility. A reserved field follows the PDE field. The reserved field may provide means for future expansion of system functions. In some embodiments packing subheaders may be used to store some header information in the payload as well; and any of the header information may be stored in the packing subheaders. In embodiments utilizing packing subheaders, one of the reserved bits would be utilized to indicate the whether or not packing subheaders are present. Such a bit might be called a packing subheader present field (PSP). Packing subheaders can be of various lengths and describe the length of the individual SDU or fragment payloads that follow each packing subheader. Alternative downlink PDU formats may be similar to the downlink PDU format **800** illustrated in FIG. **8** with minor deviations for differing characteristics.

FIG. **9** is a functional block diagram of components in an exemplary base station communications processor that process and transfer data. FIG. **9** is a high level diagram of those functional components that process and handle data packets being transferred from data source to user and vice versa. These functional components may be located in the communications processor of both the base station and the nodes. The higher layer interface **910** receives the SDUs that come to the base station from the various data sources via the backhaul interface and the input/output control.

Alternatively, the higher layer interface **910** is part of the input/output control. The higher layer interface **910** passes the SDUs to a classification module **920** that determines the connection (or destination), type and size of the SDU. This

determination is accomplished using control protocols that are unique to each particular higher layer protocol being transported. The classification data is forwarded to other base station communications processor modules to facilitate other functions such as queuing, packing, fragmentation and assigning proper PDU header characteristics. The SDUs are then forwarded to the convergence sublayer **925** for convergence subprocessing. The convergence subprocesses and their service access points provide the interfaces to higher communications protocol layers for service specific connection establishment, maintenance and data transfer. Convergence subprocesses of data are well-known in the art. One such convergence subprocess is described in a text entitled "Asynchronous Transfer Mode (ATM), Technical Overview", second edition, Harry J. R. Dutton and Peter Lenhard, published by Prentice Hall, October 1995, at pp. 3-21 through 3-24.

Upon processing by the convergence sublayer **925**, the SDUs are ready for further processing. The SDUs are distinguished by their type of message format and their connection identification information, among other things as provided by the classification module. In the data queuing module **930** and the bandwidth allocation/process bandwidth request/fragmenting/packing (BPFP) module **935** the SDUs are stored and sorted based upon their individual characteristics and various system protocols. This information may pertain to the type of user connection being served, the node the SDU is to be sent to, the type of SDU, the length of the SDU, the available physical slots in a relevant PDU, as well as many other factors. In one embodiment, the base station maps and allocates bandwidth for both the uplink and downlink communications subframes. These maps can be developed and maintained by the base station communications control modules (conveyed in FIG. **10**) in conjunction with the BPFP module **935** and may be referred to as the Uplink Subframe Maps and Downlink Subframe Maps. The communications processor must allocate sufficient bandwidth to accommodate the bandwidth requirements imposed by high priority constant bit rate (CBR) services such as Tl, El and similar constant bit rate services and their respective formats. In addition, the communications processor must allocate the remaining system bandwidth to mid-priority services and also to the lower priority services such as Internet Protocol (IP) data services and their respective formats. In one embodiment, the communications processor distributes bandwidth among these lower priority services using various QoS dependent techniques such as fair-weighted queuing and round-robin queuing, among others.

The BPFP module **935** also utilizes the data queuing module **930** to pack the SDUs into PDUs. While the SDUs are being packed into PDUs, it may be necessary to fragment an SDU if the remaining space in the relevant PDU cannot store the whole SDU. In one embodiment, fragmentation and packing occur cooperatively so as to maximize the benefit of each. For packing and fragmentation to occur in a cooperative manner, both processes should occur nearly contemporaneous to one another and in accordance with one another. If packing and fragmentation are done independently of one another, not only may the advantages of both processes be lost, the resulting system may actually be less efficient than if only one of the two processes occurred. In one embodiment the packing and fragmenting processes occur independently of the bandwidth allocation process and simply pack and fragment the SDUs as they are queued up by a separate queuing process. In another embodiment the packing and fragmentation occur in conjunction with bandwidth allocation processes and algorithms to most efficiently utilize the communications link at any one time. Numerous queuing techniques and QoS systems may be implemented, but certain embodiments should be flexible and allow the system controls to be adjusted as bandwidth demands change, as connection topography changes and as system demands change based on user requests and feedback. The variety of system configurations available and the ability to change as needed make such embodiments highly useful and largely advantageous over existing systems.

Messages now in PDU format may then be encrypted for their secure transmission. A module, such as encryption module **940**, is advantageously provided for this function. As discussed above with respect to FIG. **5**, PDU packets may then undergo a transmission convergence (TC) process to map the PDU packets into TC/PHY packets, or TDUs as previously illustrated in FIGS. **5** and **6**. This process may occur in one or multiple modules such as transmission convergence module **945**. In past systems, fragmentation may have occurred in the TC process, however, such processing at this stage would be independent of the packing and bandwidth allocation processes and could therefore result in a sharp decrease in the potential benefits provided by each of those processes. The TC process may be an intermediate format as well as a couple between the PDU formation and mapping to PHY elements (PIs as also mentioned above with respect to FIGS. **5** and **6**) in the physical (PHY) layer. Beyond the TC process, the information may then be transferred to the physical layer for further processing by the modem and the antenna, if such a transmission mechanism is utilized. This mapping of PDUs to PSs may occur in the transmit to PHY module **950** or any functional equivalent.

On the uplink processing of information, data is received from the transmission mechanism and processed by the modem. The PHY reception module **955** in the communications processor then receives the data in the PHY layer. The data undergoes an uplink transmission convergence, as described above with respect to FIG. **6**, which converts the TDU format to the encrypted PDU format. The data may then be decrypted in a decryption module **960** and passed to the BPFP module **935** for transformation from PDU packet format back to the various SDU packet formats that were originally received by the nodes from the users. The packets then undergo the convergence process in a convergence sublayer **925** in preparation for transport to the input/output control and on to the appropriate data source via the backhaul interface. By use of these exemplary data handling functional modules, data is efficiently transferred from user to data source and vice versa.

FIG. **10** is functional block diagram of components that handle control functions in an exemplary base station communications processor. The base station communications controls for an ATDD embodiment are illustrated in FIG. **10** and contain information and process control functions for each of the nodes, all of the individual communications links, and all of the system functions in order to effectively monitor, operate and optimize the communications system performance. The functional modules illustrated comprise only a high level description of exemplary control functions and there are many other functions that may occur in the control of a communications system that, for brevity, are not described herein. However, one skilled in the art will appreciate that such functions can be used in such a system.

Communications link connection setup information, maintenance information and performance information can be collected and processed during system operation, and this is advantageously performed by connection establishment and maintenance module **1010**. In one embodiment, the number

of communications links, should change infrequently as the number of nodes operating in any one sector may rarely change, while the number of connections to those particular nodes may change more rapidly. However, some embodiments may be ideally suited for those systems whose number of links rapidly change and also for those systems that rarely add or drop nodes. The power of the transmitting signal between the base station and the node can be controlled as well to maximize both signal effectiveness and efficiency. A power control module **1020** determines the most appropriate power level at which each node should transmit communications signals on the uplink. The SNR/BER calculation module **1030** can constantly measure the quality of the signal being transmitted between the base station and the nodes to determine if acceptable signal quality levels are being maintained. If not, any number of control parameters may be changed to correct the problem. The type of forward error correction, modulation, or power level utilized for the transmission of the information, among numerous other parameters, may be changed to correct any signal deficiencies.

An adaptive burst profile management module **1080** and an ATDD management module **1060** are used to control the ratio of uplink to downlink slots in each frame **202** and provide control information to the BPFP module **935** to assist its bandwidth allocation functions. A node state management module **1040** stores and utilizes information about each node to provide input into numerous control functions, such as bandwidth allocation, QoS protocol, transmission signal optimization, connection identification, and many others. As mentioned before, these exemplary controls may be utilized in embodiments practicing the current invention, but many other functions may also exist and are not mentioned here.

FIG. **11** is a functional block diagram of components that process and transfer data in an exemplary node communications processor **32**. FIG. **11** is a high level diagram of those node communications processor functional components that process and transfer SDUs from the users to the data sources and vice versa. The higher layer interface **1110** receives the SDUs that come to the node from the various connections via the connection interface. Alternatively, the higher layer interface **1110** may be part of the connection interface. The higher layer interface **1110** transfers the information to a classification module **1120** that determines the characteristics of the SDUs to be forwarded to the base station. This determination is accomplished using control protocols unique to each particular system and provides classification data necessary to correctly transfer each SDU to its respective data source. As stated above, the classification data is forwarded to other communications processor modules to facilitate functions such as queuing, packing, fragmentation and assigning proper PDU header characteristics. The SDUs are then transferred to the convergence sublayer **1125** for convergence subprocessing. As discussed above, convergence subprocessing allows various connection types from higher level communications access points to interface with the lower layers of the communications system. The convergence subprocesses of the node are similar to the convergence subprocesses previously described with respect to the base station. After convergence, the user data is transferred to the queuing module **1130** for arrangement and storing, similar to that in the base station described above, in preparation for transfer to the bandwidth allocation/create bandwidth request/packing/fragmentation (BCPF) module **1135**. This module sorts data according to connection type and various types of priority information stored within the system to determine queuing order of the various data packages. The data packages are then sequentially fitted into PDU packets as previously discussed. Again, the data packets are advantageously packed and fragmented in a coordinated manner and in the most efficient way possible to maximize the bandwidth available from frame-to-frame.

Again, it is advantageous to incorporate the packing and fragmentation processes with the bandwidth allocation process so as to maximize the flexibility, efficiency and effectiveness of both fragmentation and packing.

The SDUs, after data queuing **1130** and conversion into PDUs by BCPF module **1135** processing, are transferred to the encryption module **1140**. The encryption module can encrypt the PDUs for secure transmission, in a similar manner as that described above for the base station. Encrypted PDUs are then transmitted to the physical layer by undergoing a TC process in transmission convergence module **1145**, similar to that described above in producing TDUs. Transmission to the physical layer is completed by physical layer module **1150**, which maps the TDUs to the PIs as before. Upon transmission to physical layer, the PIs are then ready for transmission via modern, IF/RF converter and antenna to the base station, where they are processed as previously discussed and then transferred to the appropriate data source.

When data on the downlink is transmitted from the base station to the various nodes, it is received via antenna and processed by IF/RF converter modem, if such an embodiment is utilized, and arrives at the receive from physical layer module **1155**. Here, the data is received as PIs and is converted to TDUs and then undergoes a transmission convergence **1145**, as previously discussed, converting received TDUs into encrypted PDUs. The encrypted PDUs are then passed through decryption module **1160** and then are processed by unpacking and defragmentation module **1170**, which converts them back into the SDU format they were in when they arrived at the base station from the data source. Upon conversion back into the SDU format, the SDU packets then undergo a convergence process in the convergence sublayer module **1125** for communications via the higher layer interface **1110**. In the higher layer interface **1110**, the SDUs are directed to their respective connections through the connection interface, from where they are passed on to their respective connections, or users.

FIG. **12** is functional block diagram of components that handle control functions in an exemplary node communications processor. FIG. **12** illustrates some of the communications control modules that may be used by the nodes in establishing and maintaining transmission links with the base station. As mentioned before, these exemplary communications control modules are only provided for illustrative purposes, as more or fewer may be used. It is understood that other functions may be accomplished by the node communications processor that are not included in the subsequent discussion, yet those of skill in the art understand those functions to be incorporated herein. Similar to that of the base station, a connection establishment and maintenance module **1210** may be utilized to establish the communications link between the node and the base station. The node's communications processor may include other modules that also correspond closely to the base station communications control modules and perform similar functions. These modules may include a power control module **1220**, a SNR/BER calculation module **1230**, a node state management module **1240**, and an adaptive burst profile management module **1280**, among others.

These modules can perform functions that correspond to, and are complimentary with those of the base station communications control modules described above. The power control module **1220** may utilize signals sent by the base

station to adjust the node's transmission power level as necessary to optimize the communications link's performance. The SNR/BER calculation module **1230** can interact with signals from its complimentary base station control module to monitor the performance of the communications link and to request the base station to adjust the downlink transmission characteristics as necessary to optimize that performance. The node state management module **1240** can maintain information about the node and the communications link between the node and the base station and may transfer information as necessary to update such information that is stored in the base station. An adaptive burst profile management module **1280** may also be utilized to correspond with that of the base station to adaptively change uplink burst profile. In addition, the node communications controls may also include various operational parameter control modules such as automatic frequency control (AFC) and automatic gain control (AGC) control module **1190** that can control various settings of the modem used in the communications system.

FIG. **13** is a flow chart illustrating the steps of an exemplary process for forming PDUs from incoming SDUs. The process **1300** in FIG. **13** illustrates the coordination of the fragmentation and packing processes that the communications processors may utilize in one embodiment of both the base station and the nodes. An SDU or fragment that is next in the queue is identified and obtained at state **1310**; then the process moves to decision state **1315**. At decision state **1315**, the communications processor decides whether the SDU or current fragment is larger than the available bits in the payload of the current PDU. If the current SDU or fragment is larger than the available PDU payload, the process moves to state **1320**. At state **1320**, the SDU or fragment is fragmented and mapped to fill the current PDU. The process then moves to state **1325**. At state **1325**, the remaining fragment of the current SDU is buffered and can be mapped into the next successive PDU, and the process moves to state **1330**. At state **1330**, the fragmentation control bits in the PDU header are adjusted to indicate the presence and orientation of fragments in the PDU. The process then moves to state **1335** where the PDU header is updated to incorporate information regarding the payload it carries, which may include the length and the presence or absence of a packing subheader. The process then moves to state **1340**. At state **1340**, the PDU is mapped to the physical layer and transmitted either from the node to the base station, or from the base station to the node. The process then moves to state **1345**. State **1345** is a transitional state to the next PDU so that the next PDU is formed in the communications processor. The process then moves back to state **1310** where the next SDU or fragment in the queue is mapped according to the same process **1300**.

Going back to decision state **1315**, if the current SDU or fragment is not larger than the available bits in the current PDU payload, then the process moves to state **1350**. At state **1350**, the SDU or fragment is mapped to the current PDU. The PDU header is then updated in state **1355**. The process then moves to decision state **1360**. At decision state **1360**, the communications processor determines whether there are any available bits remaining in the current PDU. If there are available bits, the process then returns to state **1310** to obtain the next SDU or fragment in the queue. If there are no more available bits in the PDU, then the process moves to decision state **1365**. At decision state **1365**, the communications processor determines whether there are fragments present in the PDU. If there are fragments present in the PDU, the process then moves to state **1370**. The fragment control bits are then adjusted to indicate the presence and orientation of those fragments. The process then moves to state **1340** for PDU transmission and then returns, as before, via state **1345** to the beginning state **1310** where the next SDU or fragment is queued up. If there were no fragments present in the present PDU at state **1365**, then the fragment control bits are adjusted to indicate such absence and the process moves from state **1365** to state **1340** for PDU transmission and then onto state **1345** to create the next PDU. Through this process, PDUs are created from SDUs via a fragmentation and coordinated packing process. By coordinating the fragmenting and packing processes that occur in the bandwidth allocation process, rather than the convergence sublayer, the advantages of packing and fragmentation are optimized and maintained whereas the efficiency gained by both processes may be lost if they were performed independently of one another.

FIG. **14** is an illustration demonstrating the relationship between the PDU header, the payload and the packing subheader. FIG. **14** illustrates the packing of multiple variable length SDUs into a single PDU. The embodiment illustrated in FIG. **14** is intended to correspond with the PDU header described previously with respect to FIG. **8**. However, only the length field and the packing subheader fields described above are germane to the current discussion. An exemplary PDU **1400** contains two main sections, a PDU header section **1402** and a PDU payload section **1420**. The PDU header section **1402** can include the various components that make up the header section described in FIG. **8**. As can be seen in FIG. **14**, the parts of the PDU header section **1402** illustrated include a length field **820** and a packing subheader present **819** field. The PDU length field **820** has a sample entry of J and a sample packing subheader present entry **819** of 1 meaning, in this case, that the length of the payload is going to be J and there are packing subheaders **1404**, **1410**, **1414**, present in the payload. Packing subheaders **1404**, **1410**, **1414**, occur in the payload of the PDU packet **1400** and they can occupy a variable number of bits depending on the type of information they contain and the lengths of the corresponding SDUs that follows the packing subheaders **1404**, **1410**, **1414**.

The packing subheaders **1404**, **1410**, **1404**, may include, among other items, a length extension item (LE) and a length item. The length extension item indicates the quantity of bits required in the subheader length field to indicate the length of the SDU that follows the subheader. The length item indicates the length of the SDU. There may be multiple variable length SDUs between the second variable length SDU **1412** and a final variable length SDU **1416**; or there may be no more SDUs between the two. The PDU header **1402** contains a length field J that comprises the entire length of the payload **1420**. That payload **1420** includes the length of the first SDU **1408** (length a), the length of the second SDU **1412** (length b), the length of the last SDU **1416** (length c) as well as the lengths of the respective packing subheaders **1404**, **1410**, **1414**, and any other SDU lengths and their subheaders that are in the payload **1420**. By this system, various lengths can be utilized and accommodated while minimizing the amount of payload **1420** bits that are utilized in the packing subheaders **1404**, **1410**, and **1440**. Because the packing subheader size can be variable depending on the type of information it contains and the length of the SDU with which it corresponds, the amount of payload lost, or that is not dedicated to carrying data, is minimized, while still allowing the PDU to contain variable length SDUs in the most efficient manner.

Through the components and functions described in the preceding paragraphs, a system and method are described that utilize packing and fragmentation in an efficient manner. In certain embodiments the packing and fragmentation processes are implemented in a cooperative manner to fully realize the benefits of both. Additionally, it is advantageous to

coordinate packing and fragmentation with bandwidth allocation so that a communications system can be most flexible and able to capitalize on the circumstances that may exist in any one communications cycle. This system also utilizes a method of packing variable length SDUs that is advantageously adaptive. Through the use of variable length packing subheaders the amount of available payload that is lost in describing the data carried within it, or cell tax, is minimized further improving the effectiveness and efficiency of the packing process. It is understood that the description was mainly made with respect to a wireless data communications system. However, as stated previously, this description applies to all packeted data communications systems and it may be advantageously utilized in any one of the previously mentioned types of communications systems.

The foregoing description details certain embodiments of the invention. It will be appreciated, however, that no matter how detailed the foregoing appears in text, the invention can be practiced in many ways. As is also stated above, it should be noted that the use of particular terminology when describing certain features or aspects of the invention should not be taken to imply that the terminology is being re-defined herein to be restricted to including any specific characteristics of the features or aspects of the invention with which that terminology is associated. The scope of the invention should therefore be construed in accordance with the appended claims and any equivalents thereof.

The invention claimed is:

1. A node for a communications system that packs and fragments variable-length service data units (SDU) for mapping into variable length protocol data units (PDU), each SDU being associated with a specified connection, the node comprising:

a communications processor configured to pack and fragment SDUs associated with a specified connection into a PDU, including

allocate bandwidth for the specified connection, based on the priority of the connection,

establish a length for the PDU based on the bandwidth allocated to the specified connection in a current frame,

pack a first SDU into a payload area of the PDU,

determine whether a second SDU is larger than a remaining payload area of the PDU,

if the second SDU is not larger than the remaining payload area of the PDU, map the second SDU to the remaining payload area of the PDU, and

if the second SDU is larger than the remaining payload area of the PDU, fragment the second SDU into at least two fragments and map the first fragment to the remaining payload area of the PDU, and

include packing sub-headers in the PDU to allow determination of the length of the SDUs and the lengths of the fragments that are mapped to the PDU.

2. A node as claimed in claim **1**, wherein the communications processor is further configured to fragment the first SDU if the length of the first SDU is larger than the payload area of the PDU.

3. A node of claim **1**, wherein the first SDU is a last fragment of an SDU.

4. A node of claim **1**, wherein a packing sub-header associated with a SDU fragment comprises a fragment sequence number.

5. A node as claimed in claim **1**, wherein a packing subheader comprises a fragmentation control field indicating whether the corresponding service data unit is a first fragment, a continuing fragment, a last fragment, or an un-fragmented SDU.

6. A node as claimed in claim **1**, wherein the communications processor is further configured to prepend a header to the PDU to indicate the length of the PDU.

7. A node as claimed in claim **6**, wherein the communications processor is further configured to identify the connection in the header.

8. A method of mapping variable-length service data units (SDU) into variable length protocol data units (PDU) in a wireless communication system, where each SDU is associated with a specified connection, the method comprising:

allocating bandwidth in a current frame to a connection established at a node, based on the priority of the connection;

establishing a length for a PDU based on the bandwidth allocated to the connection;

mapping a first SDU into a payload area of the PDU, if the first SDU is not larger than the payload area of the PDU;

determining whether a second SDU is larger than a remaining payload area of the PDU;

mapping the second SDU to the remaining payload area of the PDU if the second SDU is not larger than the remaining payload area of the PDU;

fragmenting the second SDU into at least two fragments and mapping a first fragment to the remaining payload area of the PDU if the second SDU is larger than the remaining payload area of the PDU;

prepending a header to the PDU to indicate the length of the PDU; and

providing packing sub-headers in the PDU to allow determination of the length of the SDUs and the lengths of the SDU fragments that are mapped to the PDU.

9. A method as claimed in claim **8**, wherein the first SDU is fragmented if the length of the first SDU is larger than the payload area of the PDU, and a first fragment is mapped to the payload area of the PDU.

10. A method as claimed in claim **8**, wherein the first SDU is a last fragment of an SDU.

11. A method as claimed in claim **8**, wherein the packing sub-header associated with a SDU fragment comprises a fragment sequence number.

12. A method as claimed in claim **8**, wherein a packing sub-header associated with an SDU fragment comprises a fragmentation control field indicating whether the corresponding SDU or SDU fragment mapped into the PDU is a first fragment, a continuing fragment, a last fragment, or an un-fragmented SDU.

13. A method as claimed in claim **8**, further comprising identifying the connection in the header.

14. A node for a communications system that packs and fragments variable-length service data units (SDU) into variable length protocol data units (PDU), the node comprising:

a communications processor which packs and fragments SDUs associated with a specified connection into a PDU, the communication processor being configured to

establish a length for the PDU based on bandwidth currently allocated to the specified connection in a current frame based on the priority associated with the specified connection,

pack a first SDU into the payload area of the PDU, if the first SDU is not larger than payload area of the PDU and provide a corresponding packing sub-header,

fragment the second SDU into at least two fragments if a second SDU is larger than a remaining payload area of the PDU, and

pack a first fragment of the second SDU into the remaining payload area of the PDU and provide a corresponding packing sub-header,

wherein the packing sub-headers comprise a length field specifying the length of the respective corresponding SDU or SDU fragment, and further include a fragmentation control field indicating whether the respective corresponding SDU or SDU fragment is a first fragment, a continuing fragment, a last fragment, or an un-fragmented SDU.

15. The node as claimed in claim **14**, wherein the communication processor is further configured to fragment the first SDU if the length of the first SDU is larger than the payload area of the PDU, and map a first fragment to the payload area of the PDU.

16. A method of formatting protocol data units (PDUs) from incoming variable-sized service data units (SDUs) for transmission of data carried by the PDUs over a communication channel shared by one or more nodes, comprising:

provisioning a protocol data unit (PDU), including a header and a payload area, wherein the length of the PDU is established in conjunction with the bandwidth amount allocated to the specified connection in a current frame, the bandwidth amount being established frame-by-frame based on one or more communication parameters associated with the specified connection, including the priority of the specified connection, and general system parameters;

packing and fragmenting the SDUs associated with the specified connection into the payload area of the PDU based on the current length of the payload area;

adding a packing sub-header for the SDUs and the fragments that are packed and fragmented in the PDU, to allow determination of the length of the SDUs and the fragments; and

prepending a header to the PDU that indicates the length of the PDU.

17. The method of claim **16**, wherein the length of the PDU changes as the bandwidth allocated to the specified connection changes.

18. The method of claim **16**, wherein packing and fragmenting comprises:

mapping one or more SDUs into the payload area of the PDU until a remaining area in the payload area of the PDU cannot accommodate a next SDU;

fragmenting the next SDU into a first and a second fragment, the first fragment having the length of the remaining area;

mapping the first fragment to the remaining area; and

inserting fragmentation header information to indicate the fragmentation state of the payload and to identify the first fragment as being a first fragment.

19. The method of claim **18**, further comprising:

mapping the second fragment to a next PDU if the length of the second fragment fits into the length of a payload area of the next PDU; and

inserting fragmentation control information to indicate the fragmentation state of the payload and to identify the second fragment as being one of a continuing and last fragment.

20. The method of claim **19**, further comprising:

further fragmenting the second fragment if the length of the second fragment is larger than the length of the payload area of a next PDU to obtain a third fragment having the length of a payload area of the next PDU;

mapping the third fragment to the next PDU; and

inserting fragmentation control information, to indicate the fragmentation state of the payload and to identify the third fragment.

21. A method as claimed in claim **16**, further comprising multiplexing the PDU into the current frame and transmitting the frame over the wireless communication system.

22. A method of mapping variable-length service data units (SDU) into variable length protocol data units (PDU) in a wireless communication system, where each SDU is associated with a specified connection, the method comprising:

allocating bandwidth in a current frame to a connection established at a node, based on the priority of the connection;

establishing a length for a PDU based on the bandwidth allocated to the connection;

fragmenting a first SDU if the length of the first SDU is larger than the payload area of the PDU, and mapping a first fragment to the payload area of the PDU;

mapping a last fragment of the first SDU into a subsequent PDU;

determining whether a second SDU is larger than a remaining payload area of the subsequent PDU;

mapping the second SDU to the remaining payload area of the subsequent PDU if the second SDU is not larger than the remaining payload area of the subsequent PDU;

fragmenting the second SDU into at least two fragments and mapping a first fragment to the remaining payload area of the subsequent PDU if the second SDU is larger than the remaining payload area of the subsequent PDU;

prepending a header to the subsequent PDU to indicate the length of the subsequent PDU; and

providing packing sub-headers in the subsequent PDU to allow determination of the length of the SDUs and the lengths of the SDU fragments that are mapped to the subsequent PDU.

* * * * *

US008315640B2

(12) **United States Patent**
Stanwood et al.

(10) **Patent No.: US 8,315,640 B2**
(45) **Date of Patent: *Nov. 20, 2012**

(54) **METHODS AND SYSTEMS FOR TRANSMISSION OF MULTIPLE MODULATED SIGNALS OVER WIRELESS NETWORKS**

(75) Inventors: **Kenneth L. Stanwood**, Cardiff by the Sea, CA (US); **James F. Mollenauer**, Newton, MA (US); **Israel Jay Klein**, San Diego, CA (US); **Sheldon L. Gilbert**, San Diego, CA (US)

(73) Assignee: **Wi-LAN, Inc.**, Ottawa (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/645,937**

(22) Filed: **Dec. 23, 2009**

(65) **Prior Publication Data**

US 2010/0150093 A1 Jun. 17, 2010

Related U.S. Application Data

(63) Continuation of application No. 11/170,392, filed on Jun. 29, 2005, which is a continuation of application No. 09/859,561, filed on May 16, 2001, now Pat. No. 6,956,834, which is a continuation of application No. 09/316,518, filed on May 21, 1999, now Pat. No. 6,925,068.

(51) **Int. Cl.**
H04W 72/00 (2009.01)

(52) **U.S. Cl.** **455/450**; 455/422.1; 455/434; 455/451; 455/452.1; 455/453; 370/310.2; 370/322; 370/395.21; 370/395.43; 370/412

(58) **Field of Classification Search** 455/450–453, 455/422.1, 434; 370/310.2, 322, 395.21–395.43, 370/412–421

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,949,404 A 4/1976 Fletcher et al.

(Continued)

FOREIGN PATENT DOCUMENTS

EP 0353759 2/1992

(Continued)

OTHER PUBLICATIONS

"Asynchronous Transfer Mode (ATM) Technical Overview", 2nd Edition, Prentice Hall, Oct. 1995, Chapter 3, pp. 21-25.

(Continued)

Primary Examiner — Steve D Agosta
(74) *Attorney, Agent, or Firm* — Procopio, Cory, Hargreaves & Savitch LLP

(57) **ABSTRACT**

A method and apparatus for requesting and allocating bandwidth in a broadband wireless communication system. The inventive method and apparatus includes a combination of techniques that allow a plurality of CPEs to communicate their bandwidth request messages to respective base stations. One technique includes a "polling" method whereby a base station polls CPEs individually or in groups and allocates bandwidth specifically for the purpose of allowing the CPEs to respond with bandwidth requests. The polling of the CPEs by the base station may be in response to a CPE setting a "poll-me bit" or, alternatively, it may be periodic. Another technique comprises "piggybacking" bandwidth requests on bandwidth already allocated to a CPE. In accordance with this technique, currently active CPEs request bandwidth using previously unused portions of uplink bandwidth that is already allocated to the CPE. The CPE is responsible for distributing the allocated uplink bandwidth in a manner that accommodates the services provided by the CPE. By using a combination of bandwidth allocation techniques, the present invention advantageously makes use of the efficiency benefits associated with each technique.

10 Claims, 13 Drawing Sheets

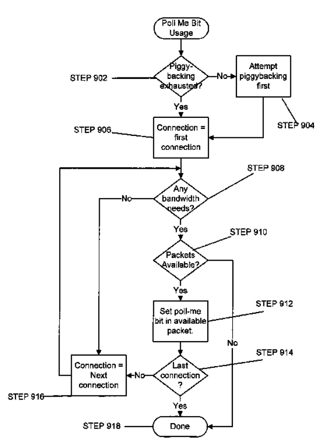

U.S. PATENT DOCUMENTS

4,464,767 A	8/1984	Bremer	
4,495,619 A	1/1985	Acampora	
4,924,461 A	5/1990	Amemiya et al.	
5,130,983 A	7/1992	Heffner, III	
5,297,144 A	3/1994	Gilbert et al.	
5,349,342 A	9/1994	Niles et al.	
5,404,374 A	4/1995	Mullins et al.	
5,420,851 A	5/1995	Seshadri et al.	
5,444,696 A	8/1995	Petranovich	
5,444,698 A	8/1995	Kito	
5,479,447 A	12/1995	Chow et al.	
5,506,848 A	4/1996	Drakopoulos et al.	
5,511,082 A	4/1996	How et al.	
5,517,503 A	5/1996	Hess	
5,592,470 A	1/1997	Rudrapatna et al.	
5,594,720 A	1/1997	Papadopoulos et al.	
5,596,577 A	1/1997	Perrault et al.	
5,615,212 A	3/1997	Ruszczyk et al.	
5,634,006 A	5/1997	Baugher et al.	
5,638,374 A	6/1997	Heath	
5,640,395 A	6/1997	Hamalainen et al.	
5,644,576 A	7/1997	Bauchot et al.	
5,673,031 A	9/1997	Meier	
5,675,573 A	10/1997	Karol et al.	
5,677,909 A	10/1997	Heide	
5,710,762 A	1/1998	Petranovich	
5,729,531 A	3/1998	Raith et al.	
5,751,708 A	5/1998	Eng et al.	
5,768,254 A	6/1998	Papadopoulos et al.	
5,787,080 A	7/1998	Hulyalkar et al.	
5,638,371 A	8/1998	Eng et al.	
5,818,820 A	10/1998	Anderson et al.	
5,818,828 A	10/1998	Packer et al.	
5,818,845 A	10/1998	Moura et al.	
5,828,695 A	10/1998	Webb	
5,859,619 A	1/1999	Wu et al.	
5,886,989 A	3/1999	Evans et al.	
5,889,963 A	3/1999	Gopal et al.	
5,890,055 A	3/1999	Chu et al.	
5,896,561 A	4/1999	Schrader et al.	
5,898,925 A	4/1999	Honkasalo et al.	
5,926,476 A	7/1999	Ghaibeh	
5,946,322 A	8/1999	Moura et al.	
5,953,344 A	9/1999	Dail et al.	
5,956,330 A	9/1999	Kerns	
5,956,338 A	9/1999	Ghaibeh	
5,956,642 A	9/1999	Larsson et al.	
5,991,287 A	11/1999	Diepstraten et al.	
6,005,853 A	12/1999	Wang et al.	
6,005,857 A	12/1999	Honkasalo et al.	
6,006,069 A	12/1999	Langston	
6,016,311 A	1/2000	Gilbert et al.	
6,016,313 A	1/2000	Foster, Jr. et al.	
6,023,458 A	2/2000	Tweedy et al.	
6,031,827 A	2/2000	Rikkinen et al.	
6,038,223 A	3/2000	Hansson et al.	
6,038,455 A	3/2000	Gardner et al.	
6,047,189 A	4/2000	Yun et al.	
6,049,549 A	4/2000	Ganz et al.	
6,049,551 A	4/2000	Hinderks et al.	
6,052,594 A	4/2000	Chuang et al.	
6,061,330 A	5/2000	Johansson	
6,075,787 A	6/2000	Bobeck et al.	
6,094,421 A	7/2000	Scott	
6,097,733 A	8/2000	Basu et al.	
6,112,080 A	8/2000	Anderson et al.	
6,114,968 A	9/2000	Ramakrishnan et al.	
6,115,390 A	9/2000	Chuah	
6,137,787 A	10/2000	Chawla et al.	
6,151,312 A	11/2000	Evans et al.	
6,212,196 B1	4/2001	Momona	
6,212,200 B1 *	4/2001	Iizuka et al.	370/468
6,222,832 B1	4/2001	Proctor	
6,249,526 B1	6/2001	Loukianov	
6,275,497 B1	8/2001	Varma et al.	
6,282,187 B1	8/2001	Evans et al.	
6,314,110 B1	11/2001	Chin et al.	
6,337,849 B1	1/2002	Smith et al.	
6,359,863 B1	3/2002	Varma et al.	
6,370,153 B1	4/2002	Eng	
6,374,112 B1	4/2002	Widegren et al.	
6,388,999 B1	5/2002	Gorsuch et al.	
6,400,684 B1	6/2002	Benmohamed et al.	
6,438,141 B1	8/2002	Hanko et al.	
6,469,991 B1	10/2002	Chuah	
6,470,016 B1	10/2002	Kalkunte et al.	
6,567,386 B1	5/2003	Sugaya et al.	
6,567,416 B1	5/2003	Chuah	
6,571,101 B1	5/2003	Schulz et al.	
6,594,246 B1	7/2003	Jorgensen	
6,615,039 B1	9/2003	Eldering	
6,628,668 B1	9/2003	Hutzli	
6,665,264 B1	12/2003	Davison et al.	
6,678,282 B2	1/2004	Sharper et al.	
6,683,866 B1	1/2004	Stanwood et al.	
6,693,891 B1	2/2004	Sugita et al.	
6,693,892 B1	2/2004	Rinne et al.	
6,731,600 B1	5/2004	Patel et al.	
6,741,614 B1	5/2004	Porter et al.	
6,771,648 B1	8/2004	Kim et al.	
6,785,252 B1	8/2004	Zimmerman et al.	
6,791,998 B1	9/2004	Yoshihara et al.	
6,795,409 B1	9/2004	Youssefmir et al.	
6,862,622 B2	3/2005	Jorgensen	
6,914,890 B1	7/2005	Tobita et al.	
6,925,068 B1	8/2005	Stanwood et al.	
6,937,566 B1	8/2005	Forslow	
6,944,148 B1	9/2005	Gehring et al.	
6,956,834 B2	10/2005	Stanwood et al.	
6,970,422 B1	11/2005	Ho et al.	
7,006,530 B2	2/2006	Stanwood	
7,050,407 B1	5/2006	Frazer et al.	
7,103,065 B1	9/2006	Quigley et al.	
7,116,682 B1	10/2006	Waclawski	
7,177,323 B2	2/2007	Yavatkar et al.	
7,418,007 B1	8/2008	Liu et al.	
7,486,639 B2	2/2009	Stanwood et al.	
7,512,154 B2	3/2009	Quigley et al.	
7,548,534 B2	6/2009	Zimmerman et al.	
7,751,437 B2	7/2010	Spinar et al.	
7,817,666 B2	10/2010	Spinar et al.	
2001/0038620 A1	11/2001	Stanwood et al.	
2002/0080816 A1	6/2002	Spinar et al.	
2004/0213197 A1	10/2004	Zimmerman et al.	
2005/0089064 A1	4/2005	Zimmerman et al.	
2006/0002336 A1	1/2006	Stanwood et al.	
2008/0232342 A1	9/2008	Spinar et al.	
2008/0232391 A1	9/2008	Spinar et al.	
2008/0253394 A1	10/2008	Spinar et al.	
2009/0168802 A1	7/2009	Spinar et al.	
2009/0175235 A1	7/2009	Spinar et al.	
2009/0207795 A1	8/2009	Spinar et al.	
2010/0150094 A1	6/2010	Stanwood et al.	
2010/0157928 A1	6/2010	Spinar et al.	
2011/0249645 A1	10/2011	Spinar et al.	
2012/0033634 A1	2/2012	Spinar et al.	

FOREIGN PATENT DOCUMENTS

EP	0507384	10/1992
EP	0622924	4/1994
EP	0720405	7/1996
EP	0845916	6/1998
EP	0891060	1/1999
EP	1045559	10/2000
GB	2320162	12/1996
WO	WO92/22162	12/1992
WO	WO9810568	3/1998
WO	WO99/38343	8/1999
WO	WO99/39532	8/1999
WO	WO00/01188	1/2000
WO	WO0072626	11/2000
WO	WO01/33772	5/2001

OTHER PUBLICATIONS

Conti et al., E-DCP, an extension of the distributed control-polling MAC protocol (DCP) for integrated services, published 1994.

Elfeitori et al., A MAC protocol for supporting real-time VBR traffic over IEEE 802.14 based HFC access networks, published May 9-12, 1999.
G. Narlikar, et al., Designing Multihop Wireless Backhand Networks with Delay Guarantees, Bell Labs, 2005, 1-4244-0222-0/06/$20.00 2006 IEEE, 12 pages.
J. Pons and J. Dunlop, "Bit Error Rate Characterisation and Modelling For GSM", IEEE 1998, pp. 3722-3727.
J.M. Torrance, L. Hanzo, "Upper Bound Performance of Adaptive Modulation in a Slow Rayleigh Fading Channel." IEEE Electronics Letters. vol. 32, Apr. 1996, p. 718.
Kim, et al., PRMA/DA: A New Media Access Control Protocol for Wireless ATM, IEEE, Jun. 23, 1996, pp. 240-244.
L.H. Charles Lee, Convolutional Coding, Fundamentals and Applications, Artech House, Inc., 1997, pp. 11-51.
Lin et al., "Error Control Coding, Fundamentals and Applications", Prentice-Hall Computer Applications in Electrical Engineering Series, 1993, pp. 315-349.
Lin, Ying-Dar. On IEEE 802.14 Medium Access Control Protocol, IEEE Communications Surveys, http://www.comsoc.org/pubs/surveys, Fourth Quarter 1998, vol. 1 No. 1, pp. 1-10.
P. Bender, et al., "CDMA/HDR: A Bandwidth Efficient High Speed Wireless Data Service for Nomadic Users", Communications Magazine, IEEE, vol. 38, No. 7, Jul. 2000, pp. 70-77.
P. Jain, "On the Impact of Channel and channel Quality Estimation on Adaptive Modulation" Dec. 2002, pp. 1-92.
Qiu et al., Dynamic Reservation Multiple Access (DRMA): A New Multiple Access Scheme for Personal Communication System (PCS), Wireless Networks 2 (1996), pp. 117-128, J.C. Baltzer AG, Science Publishers.
Redl et al., "An Introduction to GSM", Artech House, Inc., 1995; pp. 84, 85 and 95.
Sampei, et al., "Adaptive Modulation/TDMA Scheme for Personal Multi-Media Communication Systems", Proceedings of the Global Telecommunications Conference, U.S., New York, IEEE, Nov. 28, 1994, pp. 989-993.
Shannon, "A Mathematical Theory of Communication", Bell System Technical Journal, pp. 379-423 (Part 1), 623-656 (part 11), Jul. 1948.
SP-RFIv.1.1-990311 (DOCSIS1.1), dated Mar. 11, 1999, http://www.cablelabs.com/specifications/SP-RFIv1.1-I01-990311.pdf.
U.S. District Court, Northern District of California, "Amended Complaint for Declaratory Judgement Demand for Jury Trial" filed Sep. 30, 2008 in Case No. 5:08-cv-4555, 20 pages.
Ue, et al., "Symbol Rate and Modulation Level Controlled Adaptive Modulation/TDMA/TDD for Personal Communication Systems", Proceedings of the Vehicular Technology Conference, U.S., New York, IEEE, vol. Conf. 45, Jul. 25, 1995, pp. 306-310.
Wolf et al., "On the Weight Distribution of Linear Block Codes Formed From Convolutional Codes", IEEE, IEEE Transactions on Communications, vol. 44:9, Sep. 1996, pp. 1049-1051.
Ulm et al., "Data-Over-Cable Interface Specifications, Radio Frequency Interface Specification", Hewlett Packard Interim Specification Doc. Control No. SP-RFII01-970321, published Mar. 21, 1997 by MCNS Holdings, L.P., Section 6, pp. 43-85.
Mollenauer, et al., "An Efficient Media Access Control Protocol for Broadband Wireless Access Systems", Submission to IEEE 802.16 Broadband Wireless Access Working Group, Oct. 28, 1999, pp. 1-18.
Mollenauer, et al., "MAC Proposal for IEEE 802.16.1", Submission to IEEE 802.16 Broadband Wireless Access Working Group, Dec. 24, 1999, pp. 1-84.
Papadopoulos, et al., "Reduction of Mixed Co-channel Interference in Microcellular STDD Systems", Vehicular Technology Conference, 1995 IEEE 45.sup.th, vol. 2, pp. 759-763.

* cited by examiner

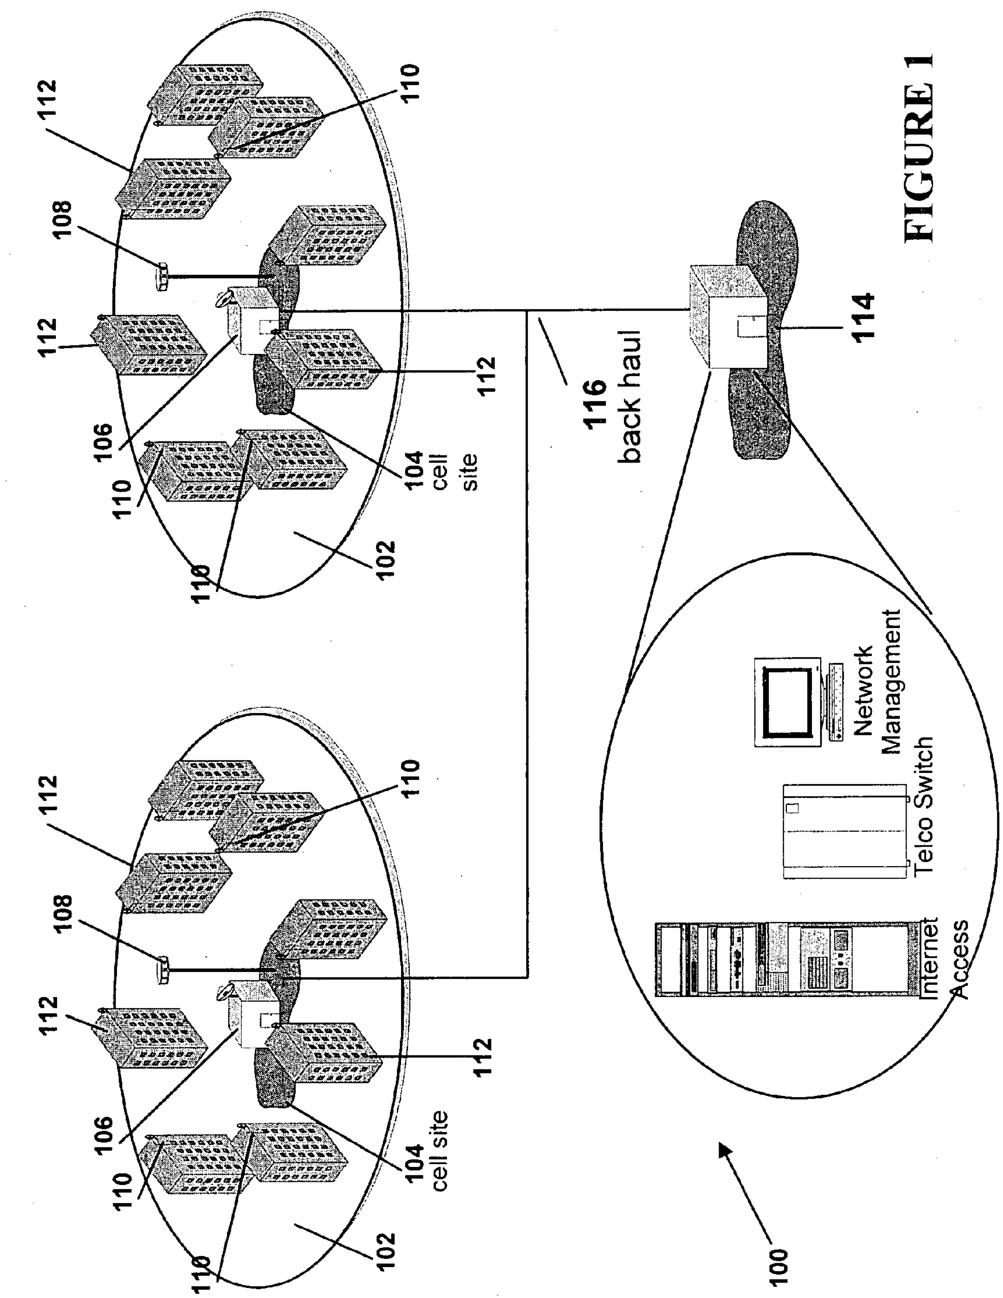

112
108
112
110
106
104 cell site
102
112
108
112
110
106
104 cell site
102
116 back haul
114
Network Management
Telco Switch
Internet Access
100

FIGURE 1

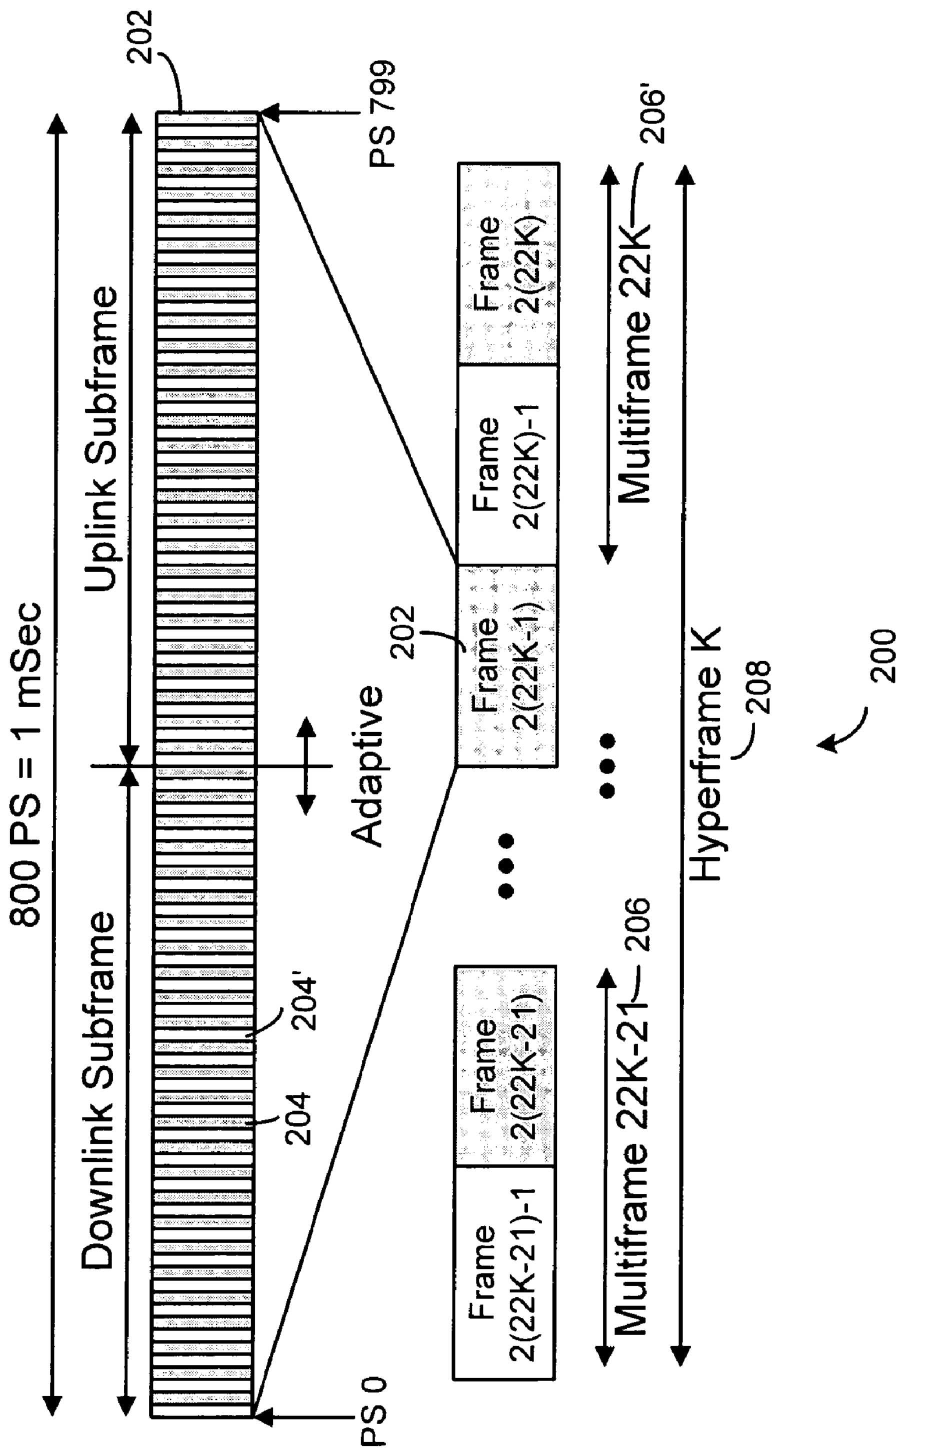
800 PS = 1 mSec
Uplink Subframe
Downlink Subframe
202
PS 799
PS 0
Adaptive
204'
204
202
Frame 2(22K-21)-1
Frame 2(22K-21)
Frame 2(22K-1)
Frame 2(22K)-1
Frame 2(22K)
Multiframe 22K-21
206
Multiframe 22K
206'
Hyperframe K
208
200

FIGURE 2

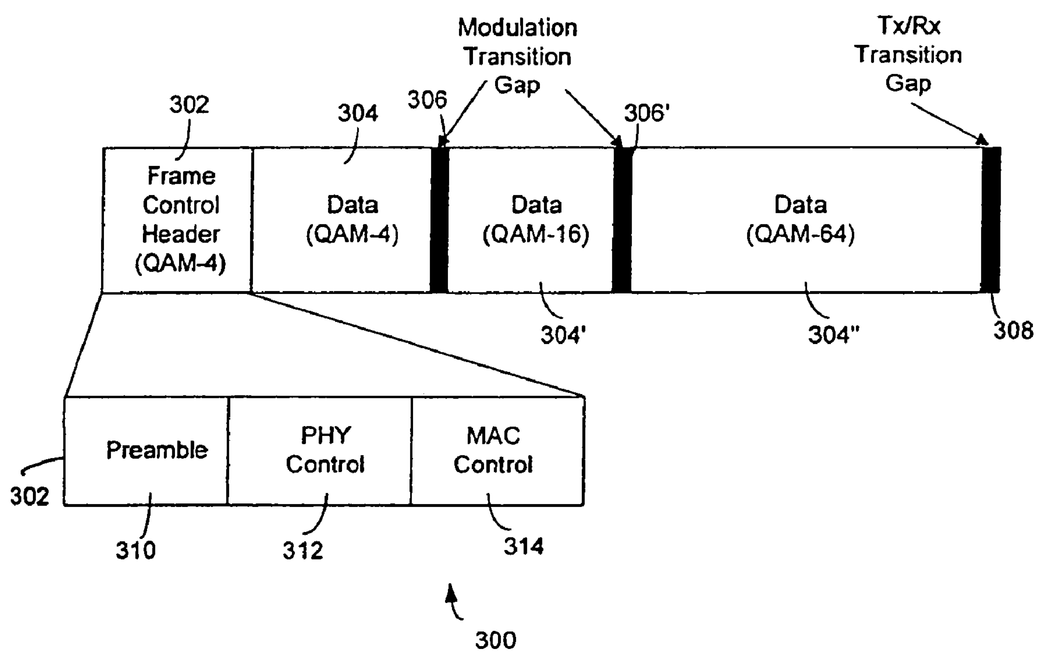
Modulation
Transition
Gap
Tx/Rx
Transition
Gap
302
304
306
306'
Frame
Control
Header
(QAM-4)
Data
(QAM-4)
Data
(QAM-16)
Data
(QAM-64)
304'
304"
308
Preamble
PHY
Control
MAC
Control
302
310
312
314
300

FIGURE 3

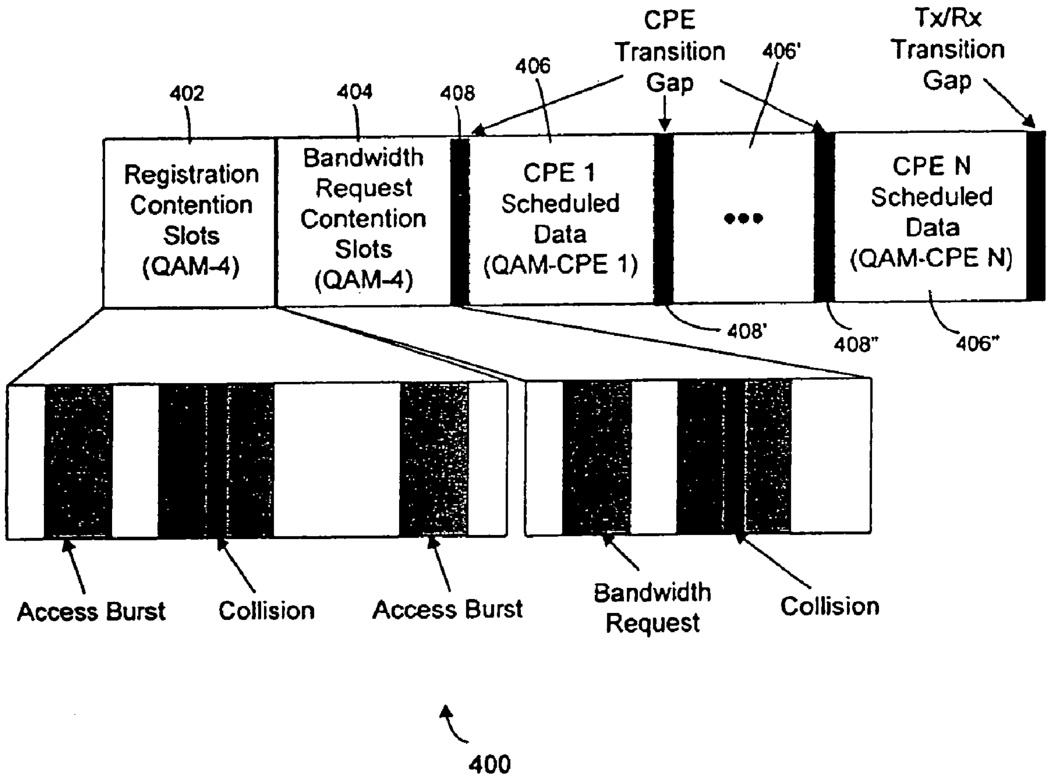

CPE
Transition
Gap
Tx/Rx
Transition
Gap
402
404
408
406
406'
Registration
Contention
Slots
(QAM-4)
Bandwidth
Request
Contention
Slots
(QAM-4)
CPE 1
Scheduled
Data
(QAM-CPE 1)
•••
CPE N
Scheduled
Data
(QAM-CPE N)
408'
408"
406"
Access Burst
Collision
Access Burst
Bandwidth
Request
Collision
400

FIGURE 4

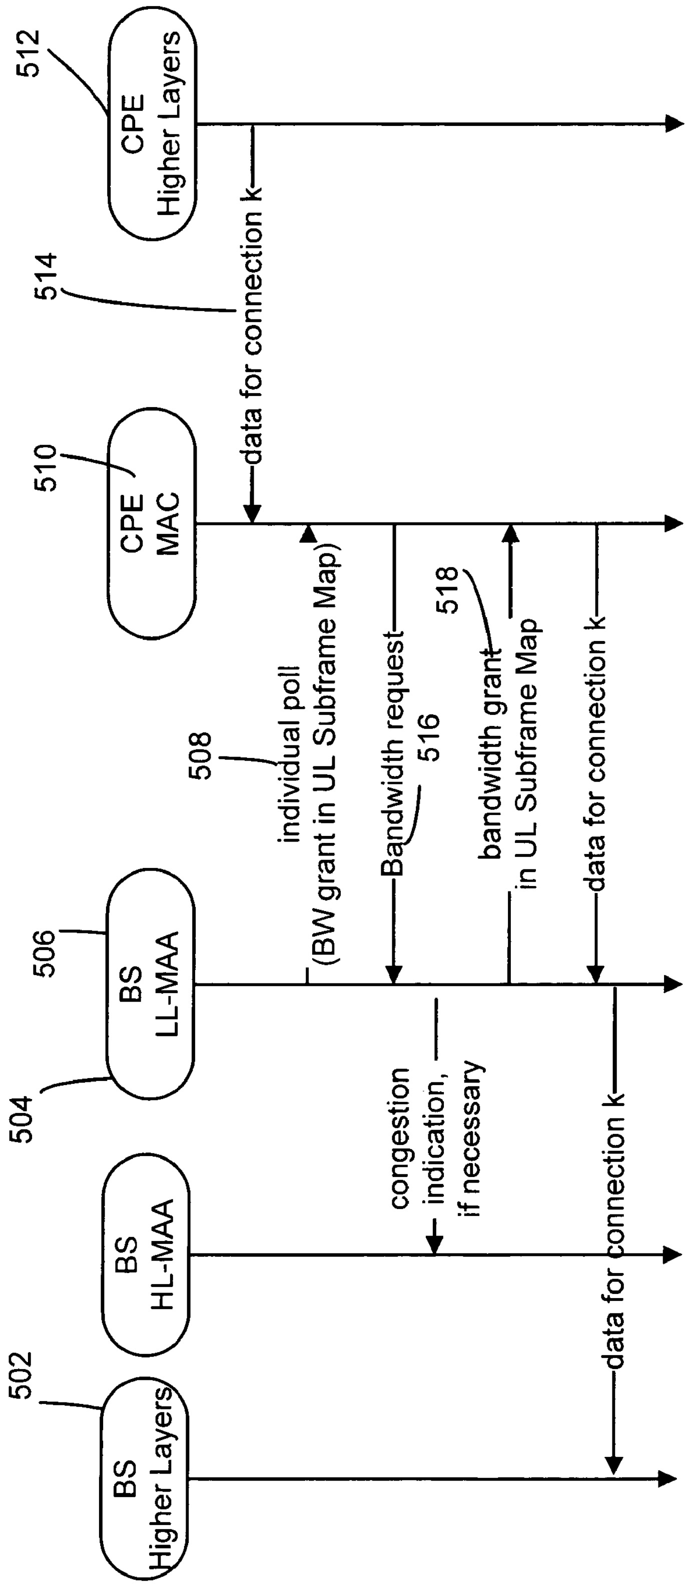
502
BS Higher Layers
BS HL-MAA
504
506
BS LL-MAA
508
individual poll (BW grant in UL Subframe Map)
510
CPE MAC
512
CPE Higher Layers
514
data for connection k
516
Bandwidth request
518
bandwidth grant in UL Subframe Map
data for connection k
congestion indication, if necessary
data for connection k

FIGURE 5

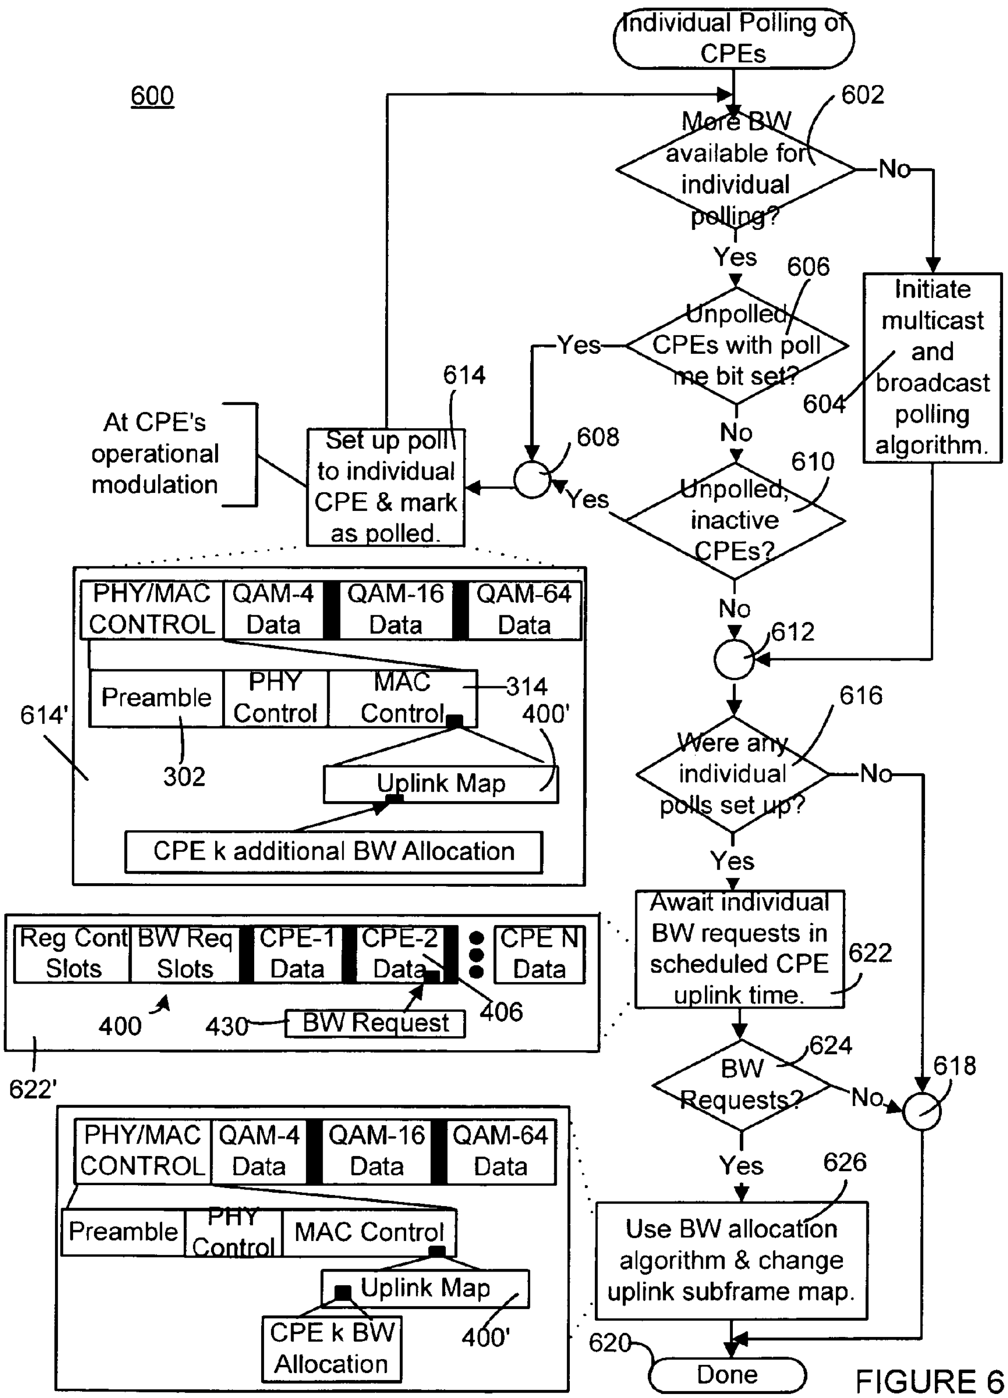
600
Individual Polling of CPEs
More BW available for individual polling?
602
No
Yes
606
Unpolled CPEs with poll me bit set?
Yes
Initiate multicast and broadcast polling algorithm.
604
No
614
At CPE's operational modulation
Set up poll to individual CPE & mark as polled.
608
Yes
610
Unpolled, inactive CPEs?
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
No
612
Preamble
PHY Control
MAC Control
314
400'
614'
302
Uplink Map
CPE k additional BW Allocation
616
Were any individual polls set up?
No
Yes
Await individual BW requests in scheduled CPE uplink time.
622
Reg Cont Slots
BW Req Slots
CPE-1 Data
CPE-2 Data
CPE N Data
400
430
BW Request
406
622'
624
BW Requests?
No
618
Yes
626
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
Uplink Map
CPE k BW Allocation
400'
Use BW allocation algorithm & change uplink subframe map.
620
Done

FIGURE 6

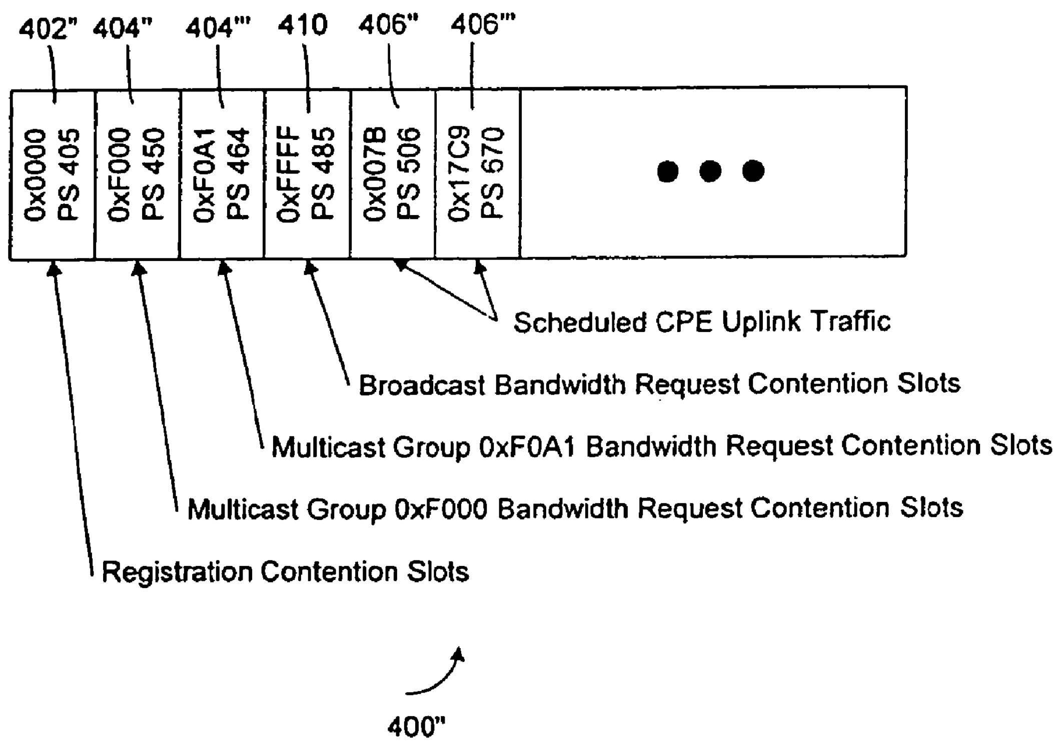

402"
404"
404'"
410
406"
406'"
0x0000 PS 405
0xF000 PS 450
0xF0A1 PS 464
0xFFFF PS 485
0x007B PS 506
0x17C9 PS 670
Scheduled CPE Uplink Traffic
Broadcast Bandwidth Request Contention Slots
Multicast Group 0xF0A1 Bandwidth Request Contention Slots
Multicast Group 0xF000 Bandwidth Request Contention Slots
Registration Contention Slots
400"

FIGURE 7

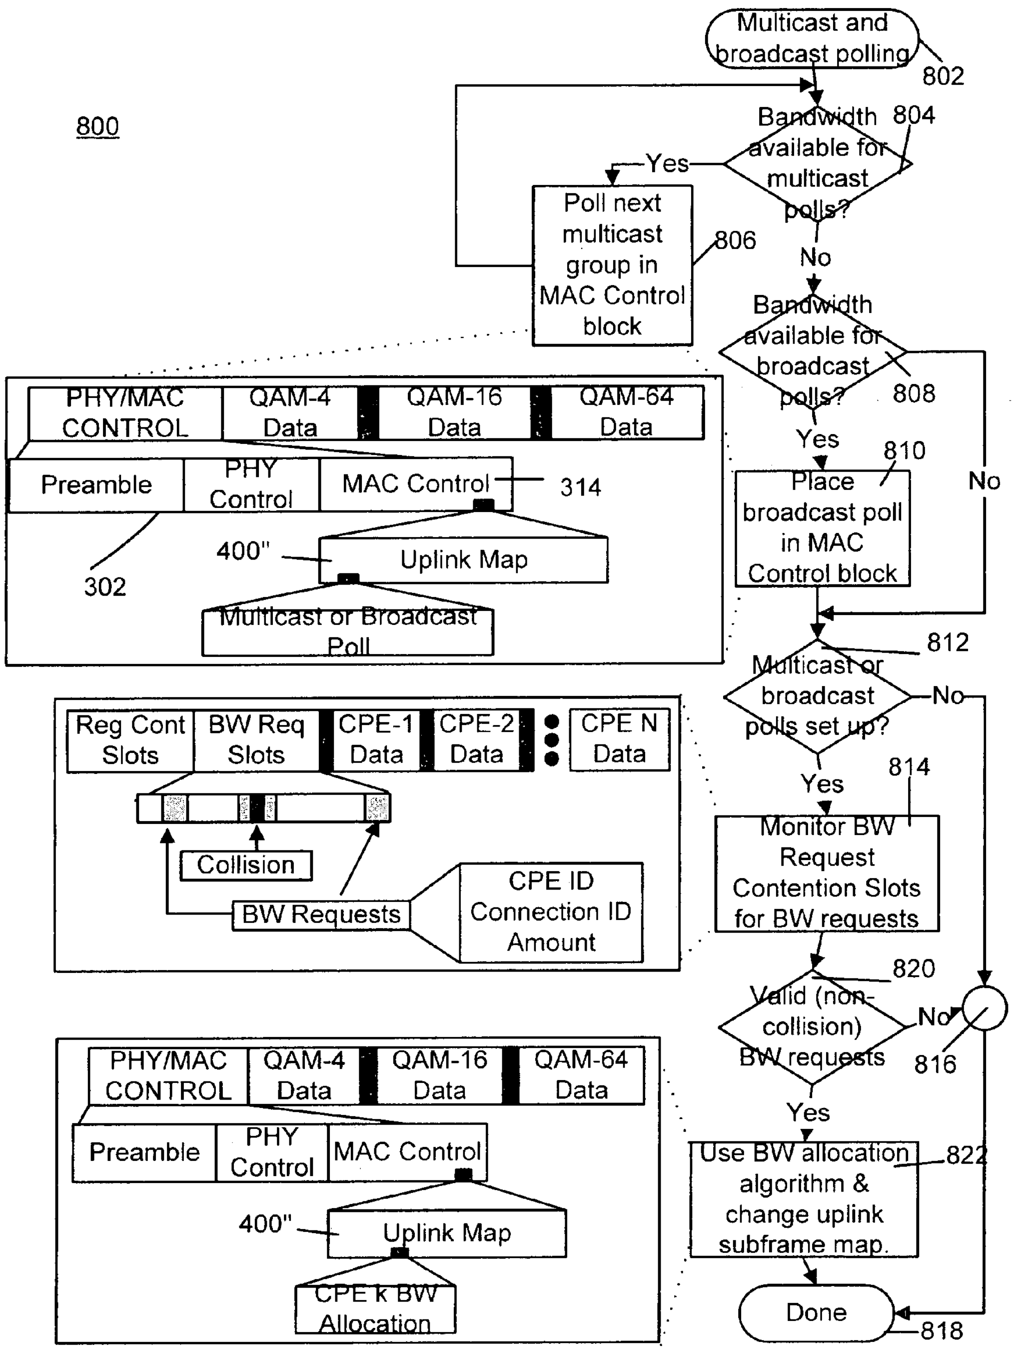
800
Multicast and broadcast polling
802
Bandwidth available for multicast polls?
804
Yes
Poll next multicast group in MAC Control block
806
No
Bandwidth available for broadcast polls?
808
Yes
Place broadcast poll in MAC Control block
810
No
Multicast or broadcast polls set up?
812
No
Yes
Monitor BW Request Contention Slots for BW requests
814
Valid (non-collision) BW requests
820
No
816
Yes
Use BW allocation algorithm & change uplink subframe map.
822
Done
818
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
314
302
400"
Uplink Map
Multicast or Broadcast Poll
Reg Cont Slots
BW Req Slots
CPE-1 Data
CPE-2 Data
CPE N Data
Collision
BW Requests
CPE ID
Connection ID
Amount
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
400"
Uplink Map
CPE k BW Allocation

FIGURE 8

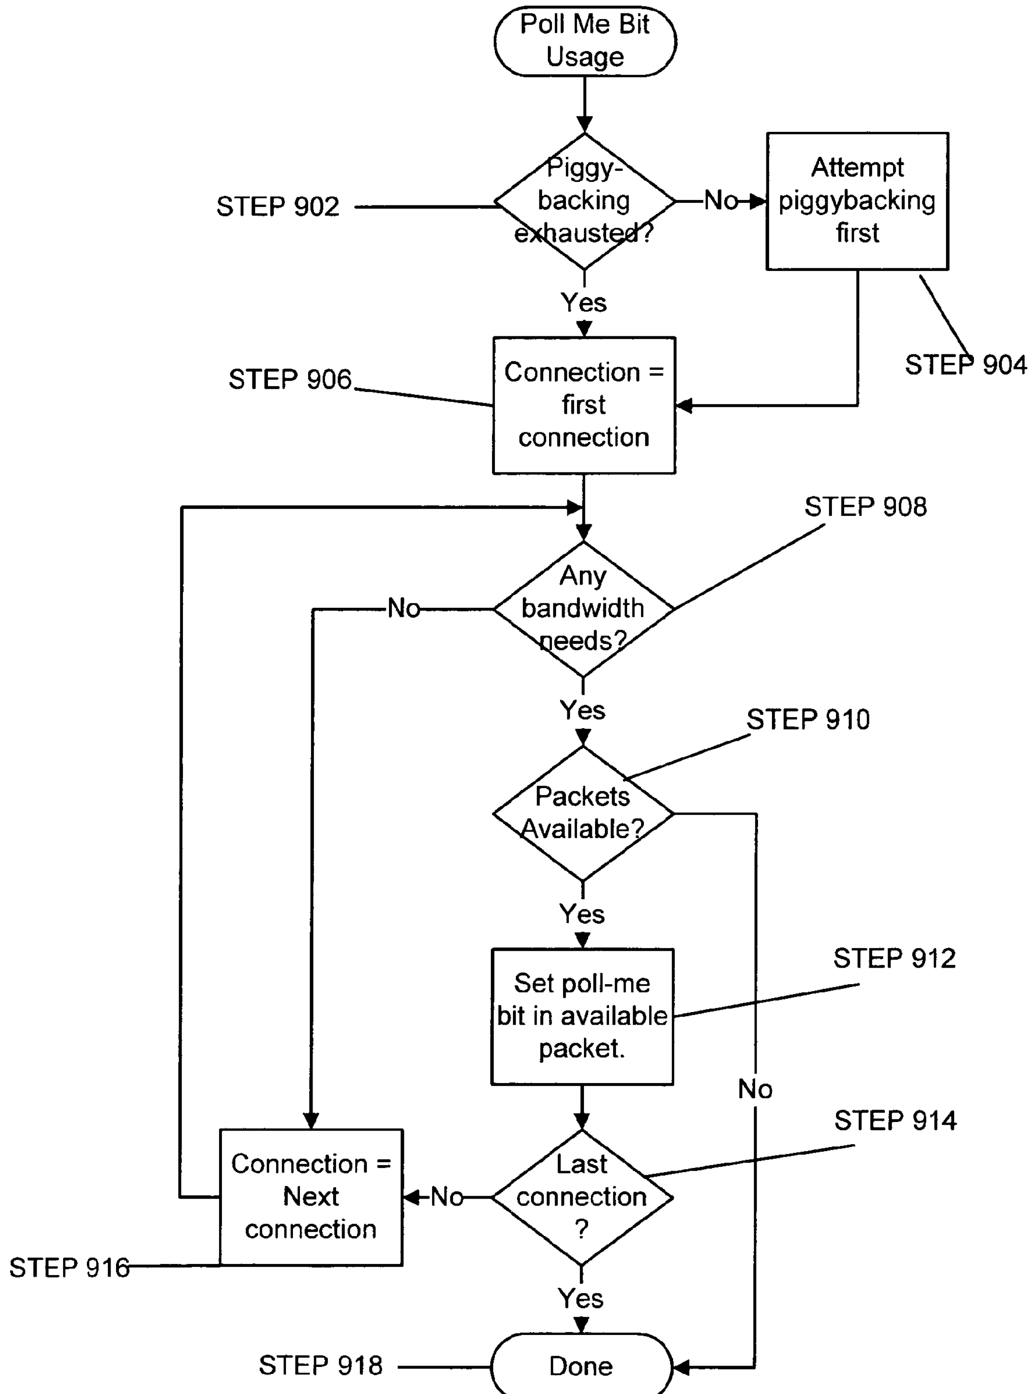
Poll Me Bit Usage
STEP 902
Piggy-backing exhausted?
No
Attempt piggybacking first
Yes
STEP 904
STEP 906
Connection = first connection
STEP 908
No
Any bandwidth needs?
Yes
STEP 910
Packets Available?
Yes
STEP 912
Set poll-me bit in available packet.
No
STEP 914
Connection = Next connection
No
Last connection ?
STEP 916
Yes
STEP 918
Done

FIGURE 9

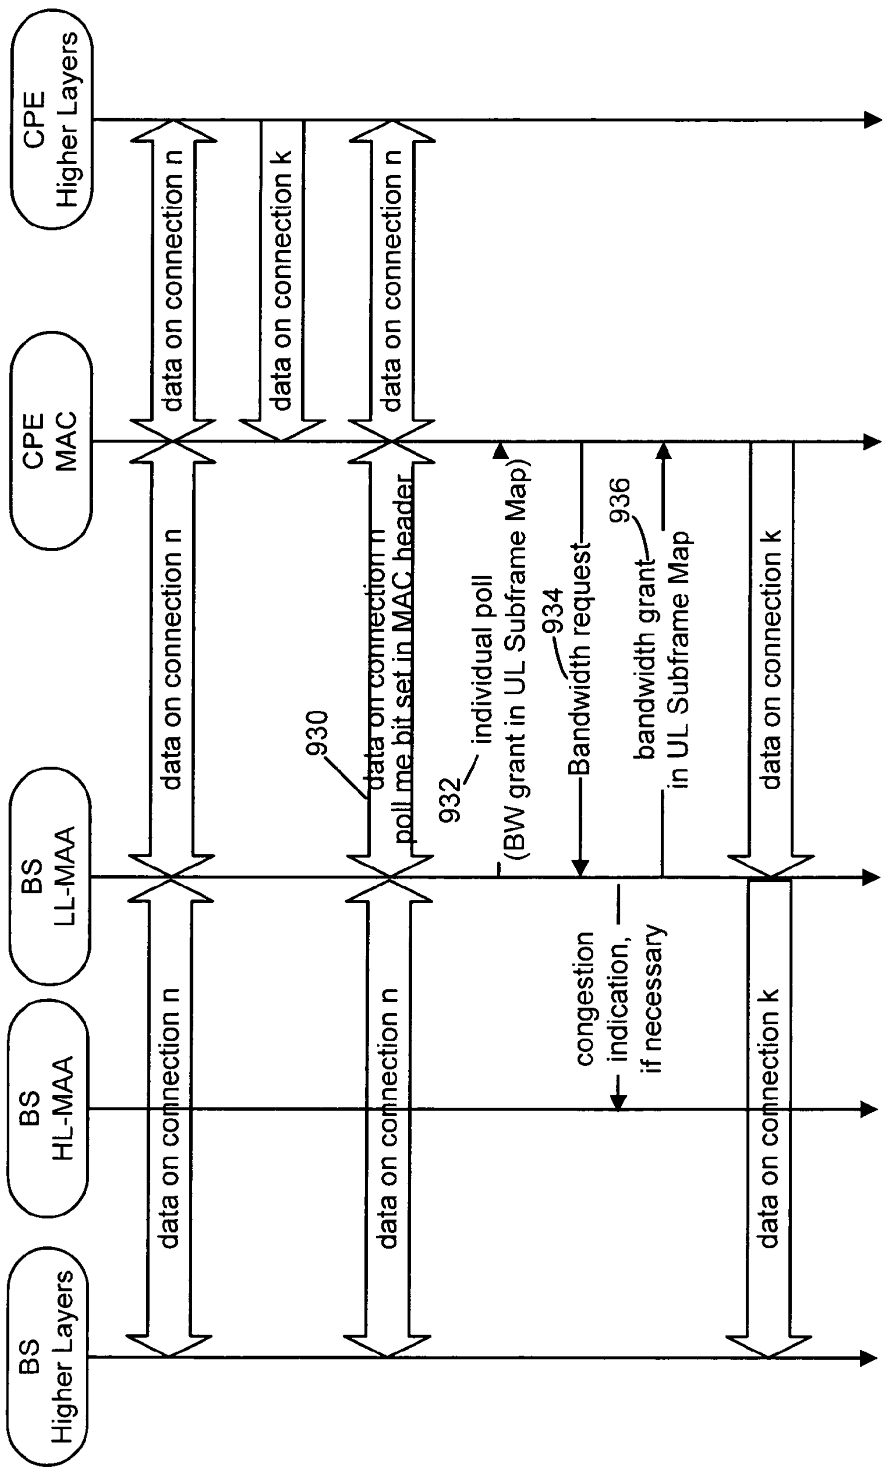
CPE Higher Layers
CPE MAC
BS LL-MAA
BS HL-MAA
BS Higher Layers
data on connection n
data on connection k
data on connection n
data on connection n
930
data on connection n poll me bit set in MAC header
932
individual poll (BW grant in UL Subframe Map)
934
Bandwidth request
936
bandwidth grant in UL Subframe Map
data on connection k
congestion indication, if necessary
data on connection n
data on connection n
data on connection k

FIGURE 10

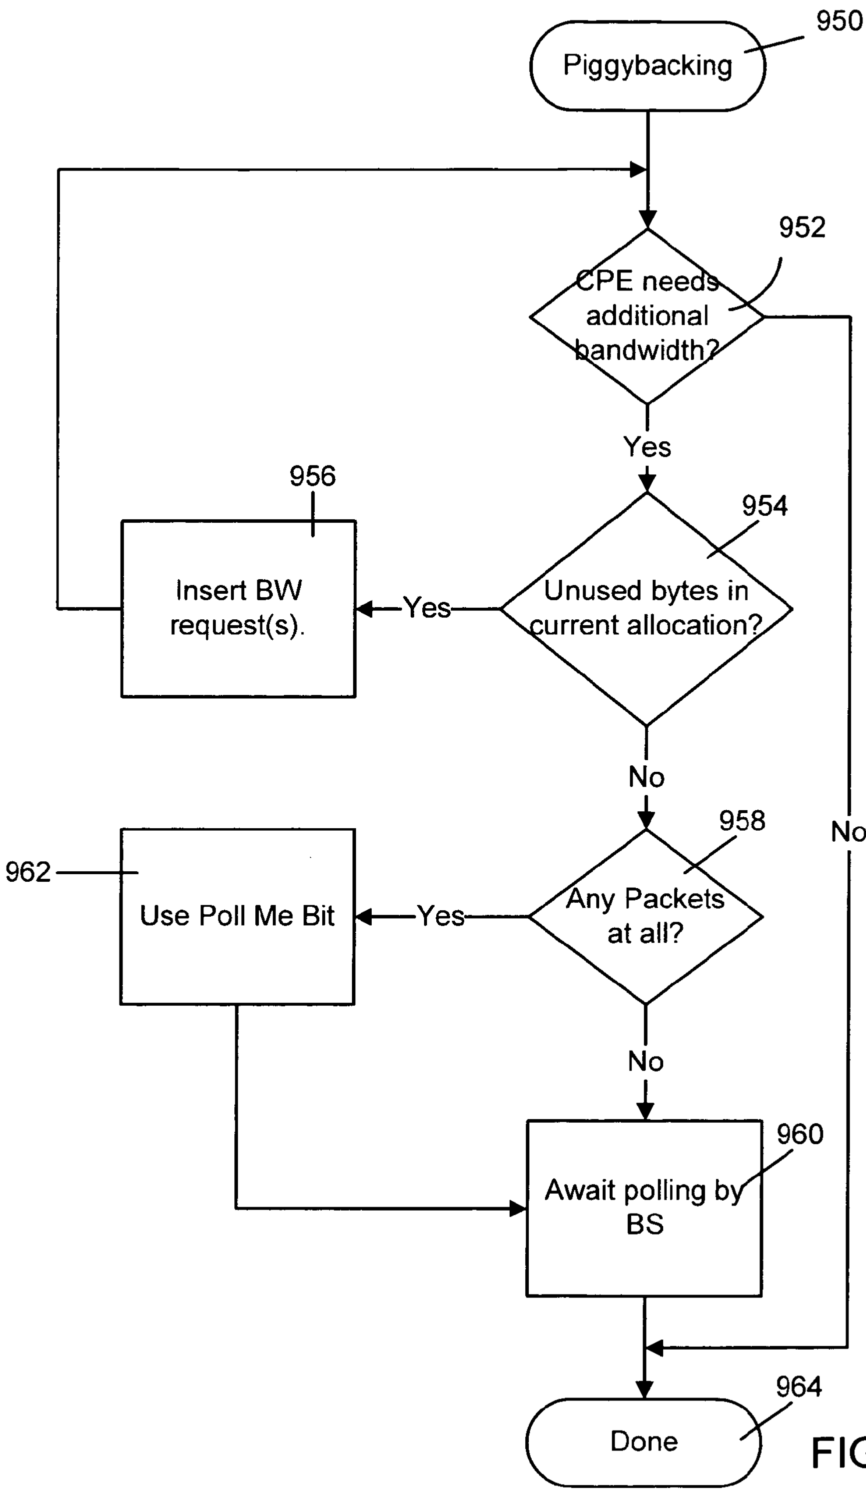

950
Piggybacking
952
CPE needs additional bandwidth?
Yes
956
Insert BW request(s).
Yes
954
Unused bytes in current allocation?
No
958
962
Use Poll Me Bit
Yes
Any Packets at all?
No
No
960
Await polling by BS
964
Done

FIGURE 11

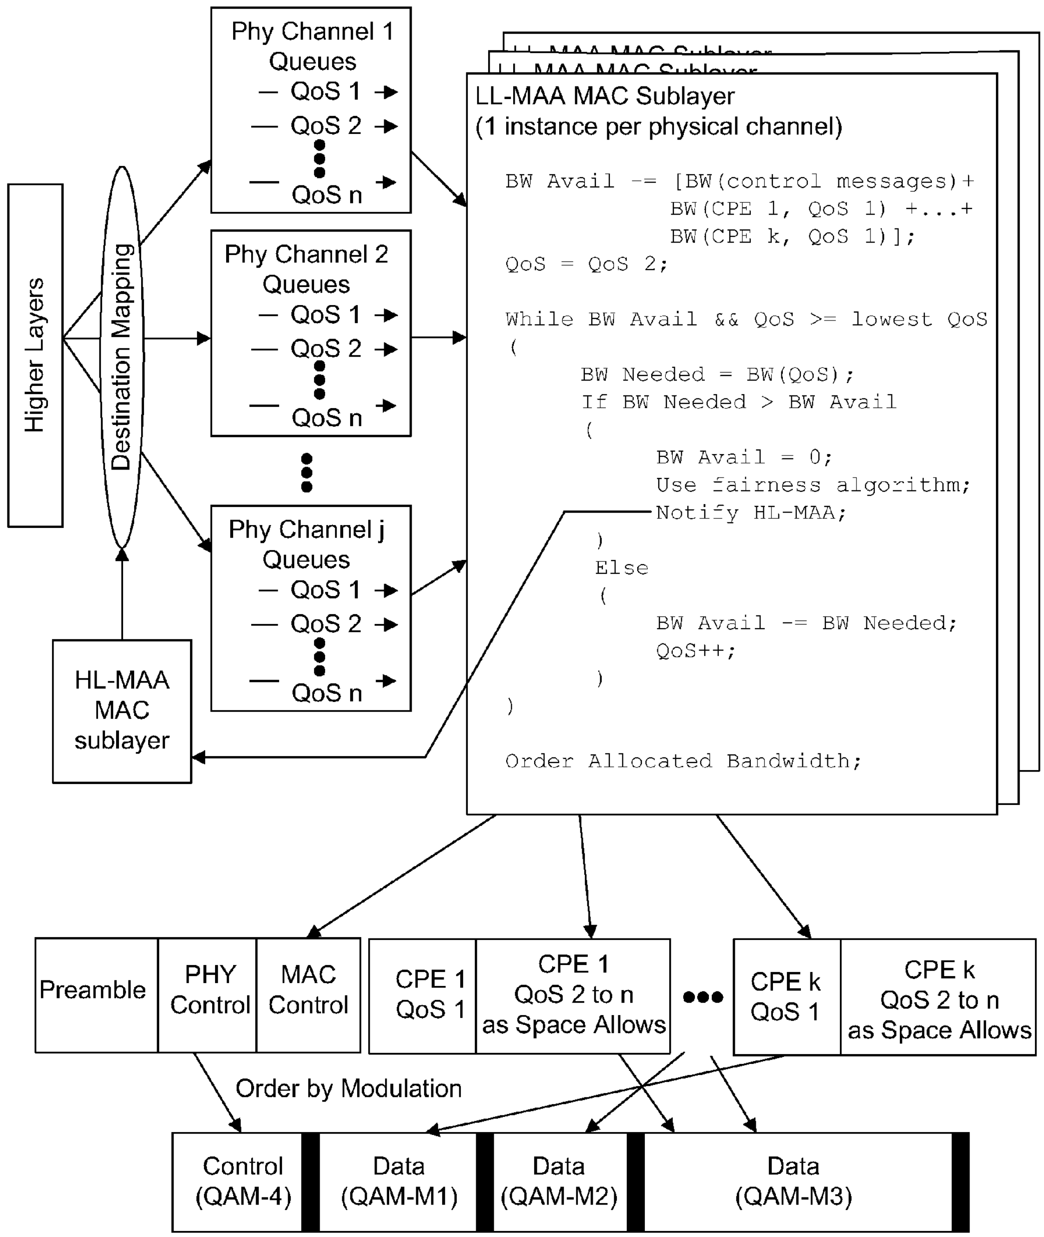
Higher Layers
Destination Mapping
Phy Channel 1 Queues
QoS 1
QoS 2
QoS n
Phy Channel 2 Queues
QoS 1
QoS 2
QoS n
Phy Channel j Queues
QoS 1
QoS 2
QoS n
HL-MAA MAC sublayer
LL-MAA MAC Sublayer
(1 instance per physical channel)
BW Avail -= [BW(control messages)+
BW(CPE 1, QoS 1) +...+
BW(CPE k, QoS 1)];
QoS = QoS 2;
While BW Avail && QoS >= lowest QoS
(
BW Needed = BW(QoS);
If BW Needed > BW Avail
(
BW Avail = 0;
Use fairness algorithm;
Notify HL-MAA;
)
Else
(
BW Avail -= BW Needed;
QoS++;
)
)
Order Allocated Bandwidth;
Preamble
PHY Control
MAC Control
CPE 1 QoS 1
CPE 1 QoS 2 to n as Space Allows
CPE k QoS 1
CPE k QoS 2 to n as Space Allows
Order by Modulation
Control (QAM-4)
Data (QAM-M1)
Data (QAM-M2)
Data (QAM-M3)

FIGURE 12

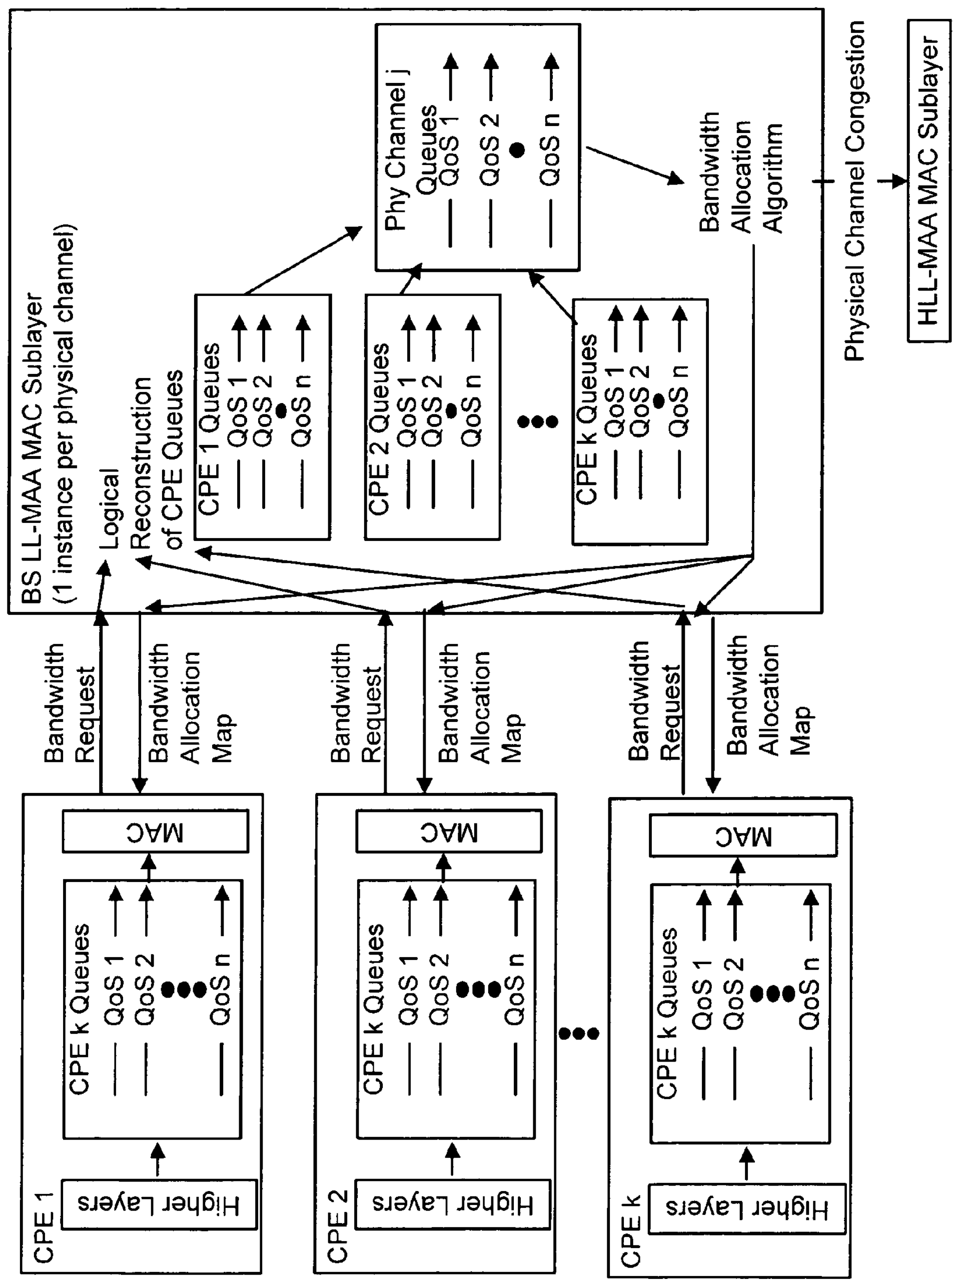
BS LL-MAA MAC Sublayer
(1 instance per physical channel)
Logical
Reconstruction
of CPE Queues
CPE 1 Queues
QoS 1
QoS 2
QoS n
CPE 2 Queues
QoS 1
QoS 2
QoS n
CPE k Queues
QoS 1
QoS 2
QoS n
Phy Channel j
Queues
QoS 1
QoS 2
QoS n
Bandwidth
Allocation
Algorithm
Physical Channel Congestion
HLL-MAA MAC Sublayer
Bandwidth
Request
Bandwidth
Allocation
Map
CPE 1
CPE 2
CPE k
MAC
CPE k Queues
QoS 1
QoS 2
QoS n
Higher Layers

FIGURE 13

METHODS AND SYSTEMS FOR TRANSMISSION OF MULTIPLE MODULATED SIGNALS OVER WIRELESS NETWORKS

CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation of U.S. patent application Ser. No. 11/170,392, filed Jun. 29, 2005, which was a continuation of U.S. patent application Ser. No. 09/859,561, filed May 16, 2001, which was a continuation of U.S. patent application Ser. No. 09/316,518, filed May 21, 1999, the disclosures of which are incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to wireless communication systems, and more particularly to a method and apparatus for efficiently allocating bandwidth between base stations and customer premises equipment in a broadband wireless communication system.

2. Description of Related Art

As described in U.S. Pat. No. 6,016,311, by Gilbert et al., issued Jan. 18, 2000, entitled "Adaptive Time Division Duplexing Method and Apparatus for Dynamic Bandwidth Allocation within a Wireless Communication System," hereby incorporated by reference, a wireless communication system facilitates two-way communication between a plurality of subscriber radio stations or subscriber units (fixed and portable) and a fixed network infrastructure. Exemplary communication systems include mobile cellular telephone systems, personal communication systems (PCS), and cordless telephones. The key objective of these wireless communication systems is to provide communication channels on demand between the plurality of subscriber units and their respective base stations in order to connect a subscriber unit user with the fixed network infrastructure (usually a wire-line system). In the wireless systems having multiple access schemes a time "frame" is used as the basic information transmission unit. Each frame is sub-divided into a plurality of time slots. Some time slots are used for control purposes and some for information transfer. Subscriber units typically communicate with the base station using a "duplexing" scheme thus allowing the exchange of information in both directions of connection.

Transmissions from the base station to the subscriber unit are commonly referred to as "downlink" transmissions. Transmissions from the subscriber unit to the base station are commonly referred to as "uplink" transmissions. Depending upon the design criteria of a given system, the prior art wireless communication systems have typically used either time division duplexing (TDD) or frequency division duplexing (FDD) methods to facilitate the exchange of information between the base station and the subscriber units. Both the TDD and FDD duplexing schemes are well known in the art.

Recently, wideband or "broadband" wireless communications networks have been proposed for providing delivery of enhanced broadband services such as voice, data and video services. The broadband wireless communication system facilitates two-way communication between a plurality of base stations and a plurality of fixed subscriber stations or Customer Premises Equipment (CPE). One exemplary broadband wireless communication system is described in the co-pending application and is shown in the block diagram of FIG. **1**. As shown in FIG. **1**, the exemplary broadband wireless communication system **100** includes a plurality of cells **102**. Each cell **102** contains an associated cell site **104** that primarily includes a base station **106** and an active antenna array **108**. Each cell **102** provides wireless connectivity between the cell's base station **106** and a plurality of customer premises equipment (CPE) **110** positioned at fixed customer sites **112** throughout the coverage area of the cell **102**. The users of the system **100** may include both residential and business customers. Consequently, the users of the system have different and varying usage and bandwidth requirement needs. Each cell may service several hundred or more residential and business CPEs.

The broadband wireless communication system **100** of FIG. **1** provides true "bandwidth-on-demand" to the plurality of CPEs **110**. CPEs **110** request bandwidth allocations from their respective base stations **106** based upon the type and quality of services requested by the customers served by the CPEs. Different broadband services have different bandwidth and latency requirements. The type and quality of services available to the customers are variable and selectable. The amount of bandwidth dedicated to a given service is determined by the information rate and the quality of service required by that service (and also taking into account bandwidth availability and other system parameters). For example, T1-type continuous data services typically require a great deal of bandwidth having well-controlled delivery latency. Until terminated, these services require constant bandwidth allocation on each frame. In contrast, certain types of data services such as Internet protocol data services (TCP/IP) are bursty, often idle (which at any one instant requires zero bandwidth), and are relatively insensitive to delay variations when active.

Due to the wide variety of CPE service requirements, and due to the large number of CPEs serviced by any one base station, the bandwidth allocation process in a broadband wireless communication system such as that shown in FIG. **1** can become burdensome and complex. This is especially true with regard to the allocation of uplink bandwidth. Base stations do not have a priori information regarding the bandwidth or quality of services that a selected CPE will require at any given time. Consequently, requests for changes to the uplink bandwidth allocation are necessarily frequent and varying. Due to this volatility in the uplink bandwidth requirements, the many CPEs serviced by a selected base station will need to frequently initiate bandwidth allocation requests. If uncontrolled, the bandwidth allocation requests will detrimentally affect system performance. If left unchecked, the bandwidth required to accommodate CPE bandwidth allocation requests will become disproportionately high in comparison with the bandwidth allocated for the transmission of substantive traffic data. Thus, the communication system bandwidth available to provide broadband services will be disadvantageously reduced.

Therefore, a need exists for a method and apparatus that can dynamically and efficiently allocate bandwidth in a broadband wireless communication system. The method and apparatus should be responsive to the needs of a particular communication link. The bandwidth needs may vary due to several factors, including the type of service provided over the link and the user type. The bandwidth allocation method and apparatus should be efficient in terms of the amount of system bandwidth consumed by the actual bandwidth request and allocation process. That is, the plurality of bandwidth requests generated by the CPE should consume a minimum percentage of available uplink bandwidth. In addition, the bandwidth allocation method and apparatus should respond

to bandwidth requests in a timely manner. Bandwidth should be allocated to high priority services in a sufficiently short time frame to maintain the quality of service specified by the CPE. Further, the bandwidth allocation method and apparatus should be capable of processing an arbitrarily large number of bandwidth allocation requests from a relatively large number of CPEs. For example, in the system shown in FIG. **1**, as many as one hundred CPEs may be allowed to be simultaneously active, coordinating their transmissions on the uplink. Furthermore, the system can accommodate approximately one thousand CPEs on the physical channel. Therefore, the need exists for a bandwidth allocation method and apparatus that can process and respond to the bandwidth allocation requests generated by a large number of CPEs.

Some prior art systems have attempted to solve bandwidth allocation requirements in a system having a shared system resource by maintaining logical queues associated with the various data sources requiring access to the shared system resource. Such a prior art system is taught by Karol et al., in U.S. Pat. No. 5,675,573, that issued on Oct. 7, 1997. More specifically, Karol et al. teach a bandwidth allocation system that allows packets or cells within traffic flows from different sources that are contending for access to a shared processing fabric to get access to that fabric in an order that is determined primarily on individual guaranteed bandwidth requirements associated with each traffic flow. In addition, the system taught by Karol et al. allow the different sources to gain access to the shared processing fabric in an order determined secondarily on overall system criteria, such as a time of arrival, or due date of packets or cells within the traffic flows. Packets or cells of data from each data source (such as a bandwidth requesting device) are queued in separate logical buffers while they await access to the processing fabric.

A need exits for a bandwidth allocation method and apparatus for efficiently processing and responding to bandwidth allocation requests. The bandwidth allocation method and apparatus should accommodate an arbitrarily large number of CPEs generating frequent and varying bandwidth allocation requests on the uplink of a wireless communication system. Such a bandwidth allocation method and apparatus should be efficient in terms of the amount of bandwidth consumed by the bandwidth request control messages exchanged between the plurality of base stations and the plurality of CPEs. In addition, the bandwidth allocation method and apparatus should respond to the bandwidth allocation requests in a timely and accurate manner. The bandwidth allocation method and apparatus should also be able to process an arbitrarily large number of bandwidth allocation requests generated by a relatively large number of CPEs. The present invention provides such a bandwidth allocation method and apparatus.

SUMMARY OF THE INVENTION

The present invention is a novel method and apparatus for requesting and allocating bandwidth in a broadband wireless communication system. The method and apparatus reduces the amount of bandwidth that must be allocated for bandwidth request and bandwidth allocation purposes. The opportunities for allowing a CPE to request bandwidth are very tightly controlled in accordance with the present invention. The present invention utilizes a combination of a number of bandwidth request and allocation techniques to control the bandwidth request process. There are a number of means by which a CPE can transmit a bandwidth request message to an associated base station.

One such means uses a "polling" technique whereby a base station polls one or more CPEs and allocates bandwidth specifically for the purpose of allowing the CPEs to respond with a bandwidth request. The polling of the CPEs by the base station may be in response to a CPE setting a "poll-me bit" or, alternatively, it may be periodic. In accordance with the present invention, periodic polls may be made to individual CPEs, to groups of CPEs, or to every CPE on a physical channel. When individually polling a CPE, the base station polls an individual CPE by allocating uplink bandwidth in an uplink sub-frame map to allow the CPE to respond with a bandwidth request. Similarly, in group polling, the base station polls several CPEs by allocating uplink bandwidth in the uplink sub-frame map to allow the CPEs to respond with a bandwidth request. The CPEs must contend for the allocated bandwidth if collisions occur. Bandwidth allocations are not in the form of an explicit message that is communicated by the base station to the CPEs, but rather the bandwidth allocations are transmitted implicitly by allocating bandwidth in the uplink sub-frame map.

Another means used by the present invention in reducing bandwidth consumed by the bandwidth request messages is the technique of "piggybacking" bandwidth requests on bandwidth already allocated to a CPE. In accordance with this technique, currently active CPEs request bandwidth using previously unused portions of uplink bandwidth that is already allocated to the CPE. Alternatively, the bandwidth requests can be piggybacked on uplink bandwidth already allocated and currently being used by a data service. In accordance with this alternative, the CPE "steals" bandwidth already allocated for a data connection by inserting bandwidth requests in time slots previously used for data.

The CPE is responsible for distributing the allocated uplink bandwidth in a manner that accommodates the services provided by the CPE. The CPE is free to use the uplink bandwidth that was allocated to it in a manner that is different than that originally requested or granted by the base station. The CPE advantageously determines which services to give bandwidth to and which services must wait for subsequent bandwidth requests. One advantage of having the CPE determine how to distribute its allocated bandwidth is that it relieves the base station from performing this task. In addition, the communication overhead that is required by having the base station instruct the CPE how to distribute its allocated bandwidth is eliminated. By using a combination of bandwidth allocation techniques, the present invention advantageously makes use of the efficiency benefits associated with each technique.

The base station media access control ("MAC") allocates available bandwidth on a physical channel on the uplink and the downlink. Within the uplink and downlink sub-frames, the base station MAC allocates the available bandwidth between the various services depending upon the priorities and rules imposed by their quality of service ("QoS"). The base station MAC maintains a set of queues for each physical channel that it serves. Within each physical channel queue set, the base station maintains a queue for each QoS. The queues hold data that is ready to be transmitted to the CPEs present on the physical channel. The base station higher MAC control layers are free to implement any convenient fairness or traffic shaping algorithms regarding the sharing of access between connections at the same QoS, without impacting the base station lower MAC control layers. In determining the amount of bandwidth to allocate at a particular QoS for a particular CPE, the base station takes into account the QoS, modulation, and the fairness criteria used to keep an individual CPE from using up all available bandwidth. In one embodiment, the

base station attempts to balance the uplink/downlink bandwidth allocations using an adaptive time-division duplexing technique (ATDD).

The uplink bandwidth allocation method is very similar to the downlink bandwidth allocation except that, rather than being maintained by the base station, the data queues are distributed across and maintained by each individual CPE. Rather than check the queue status directly, the base station preferably receives requests for bandwidth from the CPEs using the techniques described above.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. **1** shows a broadband wireless communication system adapted for use with the present invention.

FIG. **2** shows a TDD frame and multi-frame structure that can be used by the communication system of FIG. **1** in practicing the present invention.

FIG. **3** shows an example of a downlink sub-frame that can be used by the base stations to transmit information to the plurality of CPEs in the wireless communication of FIG. **1**.

FIG. **4** shows an exemplary uplink sub-frame that is adapted for use with the present bandwidth allocation invention.

FIG. **5** is a flow diagram showing the information exchange sequence used in practicing the individual polling technique of the present invention.

FIG. **6** is a flow diagram showing the individual polling technique of the present invention.

FIG. **7** shows an exemplary uplink sub-frame map that is used to facilitate the present multicast/broadcast bandwidth allocation technique.

FIG. **8** is a flow diagram showing the multicast and broadcast polling technique of the present invention.

FIG. **9** is a flow diagram showing use of a "poll-me" to stimulate polling of a CPE in accordance with the present invention.

FIG. **10** shows the message sequence that is used by the present invention in requesting polls using the "poll-me" bit.

FIG. **11** is a flow diagram showing the bandwidth request piggybacking process of the present invention.

FIG. **12** shows the downlink bandwidth allocation method used by the present invention.

FIG. **13** shows the uplink bandwidth allocation method used by the present invention.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION OF THE INVENTION

Throughout this description, the preferred embodiment and examples shown should be considered as exemplars, rather than as limitations on the present invention.

The preferred embodiment of the present invention is a method and apparatus for allocating bandwidth in a broadband wireless communication system. One very important performance criterion of a broadband wireless communication system, and any communication system for that matter having a physical communication medium shared by a plurality of users, is how efficiently the system uses the physical medium. Because wireless communication systems are shared-medium communication networks, access and transmission by subscribers to the network must be controlled. In wireless communication systems a Media Access Control ("MAC") protocol typically controls user accesses to the physical medium. The MAC determines when subscribers are allowed to transmit on the physical medium. In addition, if contentions are permitted, the MAC controls the contention process and resolves any collisions that occur.

In the system shown in FIG. **1**, the MAC executed by software present in the base stations **106** (in some embodiments, the software may execute on processors both in the base stations and the CPE) control the transmission time for all of the CPEs **110**. The base stations **106** receive requests for transmission rights and grant these requests within the time available taking into account the priorities, service types, quality of service and other factors associated with the CPEs **110**. As described above in the background of the invention, the services provided by the CPEs **110** TDM information such as voice trunks from a PBX. At the other end of the service spectrum, the CPEs may uplink bursty yet delay-tolerant computer data for communication with the well-known World Wide Web or Internet.

The base station MAC maps and allocates bandwidth for both the uplink and downlink communication links. These maps are developed and maintained by the base station and are referred to as the Uplink Sub-frame Maps and Downlink Sub-frame Maps. The MAC must allocate sufficient bandwidth to accommodate the bandwidth requirements imposed by high priority constant bit rate (CBR) services such as T1, E1 and similar constant bit rate services. In addition, the MAC must allocate the remaining system bandwidth across the lower priority services such as Internet Protocol (IP) data services. The MAC distributes bandwidth among these lower priority services using various QoS dependent techniques such as fair-weighted queuing and round-robin queuing.

The downlink of the communication system shown in FIG. **1** operates on a point-to-multi-point basis (i.e., from the base station **106** to the plurality of CPEs **110**). As described in U.S. Pat. No. 6,016,311, by Gilbert et al., issued Jan. 18, 2000, entitled "Adaptive Time Division Duplexing Method and Apparatus for Dynamic Bandwidth Allocation within a Wireless Communication System," the central base station **106** includes a sectored active antenna array **108** which is capable of simultaneously transmitting to several sectors. In one embodiment of the system **100**, the active antenna array **108** transmits to six independent sectors simultaneously. Within a given frequency channel and antenna sector, all stations receive the same transmission. The base station is the only transmitter operating in the downlink direction, hence it transmits without having to coordinate with other base stations, except for the overall time-division duplexing that divides time into upstream (uplink) and downstream (downlink) transmission periods. The base station broadcasts to all of the CPEs in a sector (and frequency). The CPEs monitor the addresses in the received messages and retain only those addressed to them.

The CPEs **110** share the uplink on a demand basis that is controlled by the base station MAC. Depending upon the class of service utilized by a CPE, the base station may issue a selected CPE continuing rights to transmit on the uplink, or the right to transmit may be granted by a base station after receipt of a request from the CPE. In addition to individually addressed messages, messages may also be sent by the base station to multicast groups (control messages and video distribution are examples of multicast applications) as well as broadcast to all CPEs.

Within each sector, in accordance with the present invention, CPEs must adhere to a transmission protocol that minimizes contention between CPEs and enables the service to be tailored to the delay and bandwidth requirements of each user application. As described below in more detail, this transmission protocol is accomplished through the use of a polling mechanism, with contention procedures used as a backup

mechanism should unusual conditions render the polling of all CPEs unfeasible in light of given delay and response-time constraints. Contention mechanisms can also be used to avoid individually polling CPEs that are inactive for long time periods. The polling techniques provided by the present inventive method and apparatus simplifies the access process and guarantees that service applications receive bandwidth allocation on a deterministic basis if required. In general, data service applications are relatively delay-tolerant. In contrast, real-time service applications such as voice and video services require that bandwidth allocations be made in a timely manner and in adherence to very tightly-controlled schedules.

Frame Maps—Uplink and Downlink Sub-Frame Mappings

In one preferred embodiment of the present invention, the base stations **106** maintain sub-frame maps of the bandwidth allocated to the uplink and downlink communication links. As described in U.S. Pat. No. 6,016,311, by Gilbert et al., issued Jan. 18, 2000, entitled "Adaptive Time Division Duplexing Method and Apparatus for Dynamic Bandwidth Allocation within a Wireless Communication System," the uplink and downlink are preferably multiplexed in a time-division duplex (or "TDD") manner. In one embodiment, a frame is defined as comprising N consecutive time periods or time slots (where N remains constant). In accordance with this "frame-based" approach, the communication system dynamically configures the first N.sub.1 time slots (where N is greater than or equal to N.sub.1) for downlink transmissions only. The remaining N.sub.2 time slots are dynamically configured for uplink transmissions only (where N.sub.2 equals N-N.sub.1). Under this TDD frame-based scheme, the downlink sub-frame is preferably transmitted first and is prefixed with information that is necessary for frame synchronization.

FIG. **2** shows a TDD frame and multi-frame structure **200** that can be used by a communication system (such as that shown in FIG. **1**) in practicing the present invention. As shown in FIG. **2**, the TDD frame is subdivided into a plurality of physical slots (PS) **204**. In the embodiment shown in FIG. **2**, the frame is one millisecond in duration and includes 800 physical slots. Alternatively, the present invention can be used with frames having longer or shorter duration and with more or fewer PSs. The available bandwidth is allocated by a base station in units of a certain pre-defined number of PSs. Some form of digital encoding, such as the well-known Reed-Solomon encoding method, is performed on the digital information over a pre-defined number of bit units referred to as information elements (PI). The modulation may vary within the frame and determines the number of PS (and therefore the amount of time) required to transmit a selected PI.

As described in U.S. Pat. No. 6,016,311, by Gilbert et al., issued Jan. 18, 2000, entitled "Adaptive Time Division Duplexing Method and Apparatus for Dynamic Bandwidth Allocation within a Wireless Communication System," in one embodiment of the broadband wireless communication system shown in FIG. **1**, the TDD framing is adaptive. That is, the number of PSs allocated to the downlink versus the uplink varies over time. The present bandwidth allocation method and apparatus can be used in both adaptive and fixed TDD systems using a frame and multi-frame structure similar to that shown in FIG. **2**. As shown in FIG. **2**, to aid periodic functions, multiple frames **202** are grouped into multi-frames **206**, and multiple multi-frames **206** are grouped into hyper-frames **208**. In one embodiment, each multi-frame **206** comprises two frames **202**, and each hyper-frame comprises twenty-two multi-frames **206**. Other frame, multi-frame and hyper-frame structures can be used with the present invention. For example, in another embodiment of the present invention, each multi-frame **206** comprises sixteen frames **202**, and each hyper-frame comprises thirty-two multi-frames **206**. Exemplary downlink and uplink sub-frames used to in practicing the present invention are shown respectively in FIGS. **3** and **4**.

Downlink Sub-Frame Map

FIG. **3** shows one example of a downlink sub-frame **300** that can be used by the base stations **106** to transmit information to the plurality of CPEs **110**. The base station preferably maintains a downlink sub-frame map that reflects the downlink bandwidth allocation. The downlink sub-frame **300** preferably comprises a frame control header **302**, a plurality of downlink data PSs **304** grouped by modulation type (e.g., PS **304** data modulated using a QAM-4 modulation scheme, PS **304'** data modulated using QAM-16, etc.) and possibly separated by associated modulation transition gaps (MTGs) **306** used to separate differently modulated data, and a transmit/receive transition gap **308**. In any selected downlink sub-frame any one or more of the differently modulated data blocks may be absent. In one embodiment, modulation transition gaps (MTGs) **306** are 0 PS in duration. As shown in FIG. **3**, the frame control header **302** contains a preamble **310** used by the physical protocol layer (or PHY) for synchronization and equalization purposes. The frame control header **302** also includes control sections for both the PHY (**312**) and the MAC (**314**).

The downlink data PSs are used for transmitting data and control messages to the CPEs **110**. This data is preferably encoded (using a Reed-Solomon encoding scheme for example) and transmitted at the current operating modulation used by the selected CPE. Data is preferably transmitted in a pre-defined modulation sequence: such as QAM-4, followed by QAM-16, followed by QAM-64. The modulation transition gaps **306** contain preambles and are used to separate the modulations. The PHY Control portion **312** of the frame control header **302** preferably contains a broadcast message indicating the identity of the PS **304** at which the modulation scheme changes. Finally, as shown in FIG. **3**, the Tx/Rx transition gap **308** separates the downlink sub-frame from the uplink sub-frame which is described in more detail below.

Uplink Sub-Frame Map

FIG. **4** shows one example of an uplink sub-frame **400** that is adapted for use with the present bandwidth allocation invention. In accordance with the present bandwidth allocation method and apparatus, the CPEs **110** (FIG. **1**) use the uplink sub-frame **400** to transmit information (including bandwidth requests) to their associated base stations **106**. As shown in FIG. **4**, there are three main classes of MAC control messages that are transmitted by the CPEs **110** during the uplink frame: (1) those that are transmitted in contention slots reserved for CPE registration (Registration Contention Slots **402**); (2) those that are transmitted in contention slots reserved for responses to multicast and broadcast polls for bandwidth allocation (Bandwidth Request Contention Slots **404**); and those that are transmitted in bandwidth specifically allocated to individual CPEs (CPE Scheduled Data Slots **406**).

The bandwidth allocated for contention slots (i.e., the contention slots **402** and **404**) is grouped together and is transmitted using a pre-determined modulation scheme. For example, in the embodiment shown in FIG. **4** the contention slots **402** and **404** are transmitted using a QAM-4 modulation. The remaining bandwidth is grouped by CPE. During its scheduled bandwidth, a CPE **110** transmits with a fixed modulation that is determined by the effects of environmental factors on transmission between that CPE **110** and its associated base station **106**. The downlink sub-frame **400** includes

a plurality of CPE transition gaps (CTGs) **408** that serve a similar function to the modulation transition gaps (MTGs) **306** described above with reference to FIG. **3**. That is, the CTGs **408** separate the transmissions from the various CPEs **110** during the uplink sub-frame. In one embodiment, the CTGs **408** are 2 physical slots in duration. A transmitting CPE preferably transmits a 1 PS preamble during the second PS of the CTG **408** thereby allowing the base station to synchronize to the new CPE **110**. Multiple CPEs **110** may transmit in the registration contention period simultaneously resulting in collisions. When a collision occurs the base station may not respond.

By using the bandwidth allocation method and apparatus of the present invention, scheduled uplink traffic data is bandwidth allocated to specific CPEs **110** for the transmission of control messages and services data. The CPE scheduled data is ordered within the uplink sub-frame **400** based upon the modulation scheme used by the CPEs **110**. In accordance with the present invention and in the manner described in detail below, bandwidth is requested by a CPE **110** and is subsequently granted by an associated base station **106**. All of the bandwidth allocated to a selected CPE within a given TDD frame (or alternatively an adaptive TDD frame, as the case may be) is grouped into a contiguous CPE scheduled data block **406**. The physical slots allocated for the CTGs **408** are included in the bandwidth allocation to a selected CPE **110** in the base station uplink sub-frame map.

In addition to the bandwidth that is allocated for the transmission of the various types of broadband services (i.e., the bandwidth allocated for the CPE scheduled data slots **406**), and the bandwidth allocated for CPE registration contention slots, bandwidth must also be allocated by the base station MAC for control messages such as requests for additional bandwidth allocations. As described in more detail below, in accordance with the present invention, CPEs **110** request changes to their bandwidth allocations by making bandwidth requests of their associated base stations **106**. The present inventive method and apparatus reduces the amount of bandwidth that must be set aside for these bandwidth allocation requests. In accordance with the present invention, the opportunities for requesting bandwidth are very tightly controlled. The present invention advantageously utilizes a combination of a number of techniques to tightly control the bandwidth request process. There are a number of means by which a CPE can transmit a bandwidth request message to its associated base station.

For example, one such means uses a "polling" technique whereby a base station polls one or more CPEs and allocates bandwidth specifically for the purpose of allowing the CPE(s) to transmit bandwidth requests. In accordance with this method, the polling of CPEs by the base station may be in response to a CPE setting a "poll-me bit" in an upstream direction or it may be periodic. In accordance with the present invention, periodic polls may be made to individual CPEs (referred to as "reservation-based" polling), to groups of CPEs ("multicast" polling), or to every CPE on a physical channel ("broadcast" polling). In reservation-based polling, the base station polls an individual CPE and then allocates uplink bandwidth to allow the CPE to respond with a bandwidth request. Similarly, in multicast and broadcast polling, the base station polls several CPEs and then allocates uplink bandwidth to allow the CPEs to respond with a bandwidth request. However, the CPEs must contend for the allocated bandwidth if collisions occur. Advantageously, neither the bandwidth polls nor the bandwidth allocations are in the form of explicit messages that are communicated by the base station to the CPEs. Rather, the bandwidth polls comprise unsolicited grants of bandwidth sufficient for transmitting bandwidth requests. Bandwidth allocations are implicit via bandwidth allocations occurring in the uplink sub-frame map. The polling techniques are described in more detail below with reference to FIGS. **4-10**.

As shown in FIG. **4**, a portion of the uplink bandwidth may periodically be allocated for these bandwidth allocation or CPE connection requests. The uplink sub-frame **400** includes a plurality of bandwidth request contention slots **404**. A CPE **110** must first be registered and achieve uplink synchronization with a base station before it is allowed to request bandwidth allocation. Therefore there is no need to allow for transmit time uncertainties in the length of the bandwidth request contention period. Consequently the bandwidth request contention period may be as small as a single PI, which, in one embodiment, at QAM-4 requires 6 PS. As with the registration requests, if a collision occurs, the base station may not respond to the CPE. If, however, the base station successfully receives a bandwidth request message from a CPE, it responds by allocating the CPE additional scheduled data **406** bandwidth in the uplink sub-frame **400**. The various polling techniques used by the present invention help to minimize the need to use the contention slots **404**. These techniques are described in more detail below.

Another means used by the present invention in reducing the bandwidth consumed by the bandwidth request messages is the technique of "piggybacking" bandwidth requests on bandwidth already allocated to a CPE. In accordance with this technique, currently active CPEs request bandwidth using previously unused portions of uplink bandwidth that is already allocated to the CPE. The necessity of polling CPEs is thereby eliminated. In an alternative embodiment of the present invention, bandwidth requests are piggybacked on uplink bandwidth allocated and actively being used by a data service. In accordance with this alternative, the CPE "steals" bandwidth already allocated for a data connection by inserting bandwidth requests in time slots previously used for data. The details of these piggybacking techniques are described in more detail below with reference to FIG. **11**.

Once a CPE is allocated bandwidth by the base station, the CPE, not the base station, is responsible for using the uplink bandwidth in a manner that can accommodate the services provided by the CPE. The CPE is free to use the uplink bandwidth that was allocated to it in a manner that is different than originally requested or granted by the base station. For example, the service requirements presented to a selected CPE can change after the selected CPE requests bandwidth from its associated base station. The CPE advantageously determines which services to give bandwidth to and which services must wait for subsequent bandwidth requests. To this end, the CPE maintains a priority list of services. Those services having higher priority (e.g., those services having high quality of service demands) will be allocated bandwidth before those services having lower priority (e.g., IP-type data services). If the CPE does not have sufficient bandwidth to meet its service requirements, the CPE will request additional bandwidth allocations by either setting its poll-me bit or by piggybacking a bandwidth allocation request.

One advantage of having the CPE determine how to distribute its allocated bandwidth is that it relieves the base station from performing this task. In addition, the communication overhead that is required by having the base station instruct the CPE how to distribute its allocated bandwidth is thereby eliminated, thus increasing usable system bandwidth. In addition, the CPE is in a much better position to respond to the varying uplink bandwidth allocation needs of high quality

of service data services. Therefore, the CPE can better accommodate the needs of these types of service requirements than can the base station.

The various techniques used by the present invention to enhance the efficiency of the bandwidth allocation request process are described in more detail below in the sub-sections that follow. Although these techniques are described in separate sub-sections, the present inventive method and apparatus advantageously uses all of the techniques in combination to reduce the bandwidth consumed by the bandwidth allocation requests.

Thus, the present invention advantageously makes use of the efficiency benefits associated with each bandwidth allocation technique. For example, although an individual polling technique is beneficial with regard to the ability to provide fast response times to bandwidth allocation requests, it is relatively inefficient with regard to the amount of bandwidth consumed by the bandwidth allocation process. In contrast, the group polling method is relatively efficient with regard to the bandwidth consumed by the bandwidth allocation process, but it is less efficient with regard to the ability to respond to bandwidth allocation requests. Use of a "poll-me" bit is relatively efficient when considered from both the bandwidth consumption and response time perspectives. In addition, the piggybacking technique further enhances bandwidth consumption efficiency by using previously unused portions of the bandwidth to send the bandwidth allocation requests. In contrast to the prior art approaches, the present invention advantageously uses all of these bandwidth allocation techniques in combination to maximize efficiency.

Polling

In one embodiment of the broadband wireless system **100** of FIG. **1** designed for use with the present invention, a CPE **110** is assigned a dedicated connection identifier (ID) when the CPE **110** first registers with the system **100**. The ID is used when the base station **106** exchanges control messages with the plurality of CPEs **110**. As described above, variations in bandwidth requirements (i.e., increases or decreases to bandwidth requirements) are necessary for all services transported by the system **100** with the exception of uncompressible constant bit rate, or continuous grant (CG) services. The bandwidth requirements of uncompressible CG services do not change between connection establishment and termination. The requirements of compressible CG services, such as channelized-T1 services, may increase or decrease depending on traffic.

In contrast, many of the data services facilitated by the system **100** of FIG. **1** are bursty and delay-tolerant. Because bandwidth is provided to these services on a demand assignment basis as needed these services are commonly referred to as Demand-Assigned Multiple Access or "DAMA" services. When a CPE **110** needs to request bandwidth for a DAMA service it transmits a bandwidth request message to the base station **106**. The bandwidth request messages communicate the immediate bandwidth requirements for the DAMA service. The bandwidth requirements can and typically do vary over time. The quality of service or "QoS" for the DAMA connection is established when the CPE connection is initially established with the base station. Therefore, the base station has the ability to access or "look-up" the QoS for any DAMA service that it is currently accommodating.

As described above, in accordance with the present invention, the CPEs **110** have a number of different techniques available to them for communicating bandwidth request messages to their associated base stations. One such technique is by transmitting a bandwidth request message in response to being polled by a base station. In accordance with the polling technique taught by the present invention, the base station allocates bandwidth to selected CPEs specifically for the purpose of making bandwidth requests. The bandwidth allocations may be to individual CPEs or to groups of CPEs. As described in more detail below in the sub-section that describes the group polling technique, allocations to groups of CPEs define bandwidth request contention slots that are used in resolving bandwidth request collisions. Advantageously, the bandwidth allocations are not made in the form of explicit messages, but rather they are made in the form of bandwidth allocation increases in the transmitted map describing the uplink sub-frame **400** (FIG. **4**). Polling is performed on a per-CPE basis, bandwidth is requested on a per-connection-ID basis, and bandwidth is allocated on a per-CPE basis. These concepts are described in more detail below.

Reservation-Based Polling Technique (Individual Polling)

In accordance with the present inventive method and apparatus, when a CPE is polled individually, no explicit message is transmitted to poll the selected CPE. Rather, the CPE is allocated bandwidth in the uplink sub-frame map that is sufficient to allow the CPE to respond with the bandwidth request. Specifically, the base station allocates bandwidth in the CPE scheduled data block **406** (FIG. **4**) for the selected CPE that is sufficient to allow the selected CPE to respond with a bandwidth request message. If the selected CPE does not require more bandwidth, it returns a request for zero bytes. A zero byte request (rather than no request) is used in the individual polling process because explicit bandwidth for a reply is allocated.

In accordance with the present invention, only inactive CPEs and active CPEs that explicitly request to be polled are eligible for individual polling. Active CPEs that do not set their respective "poll-me" bits in the MAC packet header will not be polled individually. These restrictions are imposed upon the bandwidth request process by the present invention and they advantageously save bandwidth compared with polling all of the CPEs individually. In one embodiment of the present invention, active CPEs respond to polling using the modulation scheme currently in use. However, inactive CPEs may respond using a QAM-4 or similarly robust modulation scheme to ensure that their transmission is sufficiently robust to be detected by the base station even under adverse environmental conditions.

The present invention advantageously ensures timely responses to requests for more bandwidth for a constant bit rate service such as a channelized T1 service in which channels may be added or dropped dynamically. To ensure that the base station responds quickly to requests for more bandwidth for a constant bit rate service, the uplink bandwidth allocated to a constant bit rate service that is not currently operating at a maximum rate is made sufficiently large to accommodate the service's current rate and a bandwidth request.

The information exchange sequence for individual polling is shown in the flow diagram of FIG. **5**. As shown in FIG. **5**, the base station preferably has several layers of control mechanisms or protocol stacks **502**, **504** and **506** that control, among other things, the bandwidth request and allocation process. The base station MAC is sub-divided into two sub-domains: (1) the HL-MAA MAC domain **504** and the LL-MAA Mac domain **506**. The LL-MAA MAC domain spans exactly a physical channel. Each physical channel requires an instance of the LL-MAA MAC domain. The HL-MAA MAC domain spans multiple physical channels, typically all in the same sector. A MAC domain comprises an HL-MAA MAC domain and the LL-MAA MAC domains associated with the physical channels within the HL-MAA MAC domain.

As shown in FIG. **5**, the base station individually polls (as indicated by control arrow **508**) a CPE by allocating bandwidth sufficient for the CPE to respond with a bandwidth request message. This bandwidth is allocated in the uplink sub-frame **400**. If the CPE MAC **510** determines that there is data to be sent for a selected connection k (typically determined by being instructed by a higher CPE control layer **512** via a control path **514**), then the CPE MAC control mechanism issues a bandwidth request **516** to the base station MAC **506**. If there is insufficient bandwidth available to the CPE **110** as determined by the base station's LL-MAA **506**, the bandwidth request will not be granted. Else, the bandwidth request will be granted and this will be implicitly communicated to the CPE MAC **510** by the base station allocating additional bandwidth to the CPE in the uplink sub-frame **400**. This is shown in FIG. **5** via the control path **518**. The CPE will then begin transmitting data to the base station over the uplink using the bandwidth that has been allocated to it.

FIG. **6** is a flow diagram showing the individual polling technique **600** provided by the present invention. As shown in FIG. **6**, the method starts at decision STEP **602** to determine whether bandwidth is available for the purpose of individually polling the CPEs. If no more bandwidth is available for individually polling the CPEs **110** then the method proceeds to STEP **604** and initiates a multicast or broadcast polling method. This multicast and broadcast polling method is described in detail in the sub-section below. However, if sufficient bandwidth is available for the purpose of individually polling CPEs, the method proceeds to a decision STEP **606** whereat a determination is made whether there are any un-polled active CPEs that have a "poll-me" bit set. If so, the method proceeds to a control point **608**. If not, the method proceeds to a decision STEP **610** whereat it determines whether there are any un-polled inactive CPEs present. If so, the method proceeds to the control point **608**. If not, the method proceeds to a control point **612**.

The present inventive method proceeds from the control point **608** to STEP **614** to individually poll the selected CPE. Thus, the method ensures that only un-polled active CPEs requesting more bandwidth (by setting their respective "poll-me" bits) and inactive CPEs are individually polled. This reduces bandwidth as compared with a polling method that would individually poll all CPEs.

As shown in FIG. **6**, at STEP **614** the base station initiates the polling of the selected CPE and marks the CPE as polled. This is shown diagrammatically in FIG. **6** in the caption box **614'**. The caption box **614'** of FIG. **6** shows the downlink sub-frame map **300** described above in FIG. **3**. The MAC control portion **314** of the MAC frame control header **302** preferably includes an uplink sub-frame map **400'**. The uplink sub-frame map **400'** is communicated to the CPE MAC when the base station transmits this information to the CPE via the downlink. As shown in FIG. **6**, and responsive to the polling STEP **614**, the base station MAC allocates additional bandwidth to the selected CPE (in FIG. **6** this CPE is referred to as CPE "k") in the uplink. This increased bandwidth allocation is communicated to the CPE k via the uplink sub-frame map **400'**. Thus, no additional bandwidth is needed to respond to the need to poll the selected CPE.

As shown in FIG. **6**, the method then returns to the decision STEP **602** to determine whether there is more bandwidth available for individually polling the CPEs. When it is determined (at the decision STEPS **606** and **610**, respectively) that there are no active CPEs having a poll-me bit set and that there are no un-polled inactive CPEs present, the method proceeds to a decision STEP **616**. At the decision STEP **616**, the method determines whether any individual polls were performed. If not, the method proceeds to a control point **618** and the method subsequently terminates at the termination step **620**. However, if individual polls were performed, the method proceeds to a STEP **622** to await the individual bandwidth requests from the CPE that was polled (e.g., CPE "k"). As shown in the caption **622'** of FIG. **6**, this bandwidth request **430** is generated by the polled CPE (e.g., CPE "k") during the CPE scheduled data block **406** scheduled for the selected CPE in the uplink sub-frame **400**. In one embodiment, all data includes a header that indicates the type of data being transmitted. For example, in this embodiment, control messages have associated CPE-unique connection identifiers that are assigned to them when the CPE registers. The structure of the control messages allows a base station to determine that a control message is a bandwidth request.

As shown in FIG. **6**, the method proceeds from STEP **622** to a decision STEP **624** to determine whether any bandwidth requests were received. If not, the method terminates. However, if so, the method proceeds to a STEP **626** whereat a bandwidth allocation method is initiated. As described in more detail below, the base station uses a preferred bandwidth allocation method to allocate bandwidth to the requesting CPE. The bandwidth allocation is indicated to the CPE by making appropriate changes to the uplink sub-frame map **400'**. The method then terminates at STEP **620**.

Contention-Based Polling Technique (Multicast and Broadcast Polling)

As described above with reference to STEP **604** of the individual polling method of FIG. **6**, if there is not sufficient bandwidth available for the purpose of individually polling the CPEs, the present invention may be used to poll the CPEs in multicast groups and a broadcast poll may be issued by the base station. Also, if there are more inactive CPEs than there is bandwidth available to individually poll them, some CPEs may be polled in multicast groups and a broadcast poll may be issued.

In accordance with one embodiment of the invention, the addressing of CPEs is preferably performed as follows: each CPE is assigned a unique permanent address (e.g., in one embodiment the CPE has a 48-bit address) that is used in the registration process; and each CPE is also given a basic connection ID (e.g., in one embodiment the CPE is given a 16-bit basic connection ID and a 16-bit control connection ID during the registration process). Each service that is provisioned for a selected CPE is also assigned a connection ID. Connection IDs are generated by the base station MAC (specifically, by the base station HL-MAA) and are unique across an HL-MAA MAC domain. The basic connection ID that is assigned when the CPE is registered with a base station is used by the base station MAC and the CPE MAC to exchange MAC control messages between the CPE and the base station. The control connection ID (also assigned during registration) is used by the base station and the CPE to exchange control and configuration information between the base station and the CPE higher levels of control.

In accordance with one embodiment of the present invention, certain connection IDs are reserved for multicast groups and broadcast messages. Of all of the addresses available a portion of them are preferably reserved for multicast use. For example, in one embodiment of the present invention, if the four most-significant bits of the connection ID are set to logical ones (hex "Fxxxx") the address is interpreted as being set aside for multicast use. In this embodiment, a total of 4K distinct multicast addresses are available. One example of such a multicast use is for the distribution of a video service.

In one preferred embodiment, the connection ID used to indicate a broadcast to all stations is (0xFFFF) (i.e., all 16 bits are set to a logical one).

Similar to the individual polling technique described above with reference to FIGS. **5** and **6**, the multicast polling message is not explicitly transmitted by the base station to the CPE. Rather, the multicast poll message is implicitly transmitted to the CPE when the base station allocates bandwidth in the uplink sub-frame map. However, rather than associating allocated bandwidth with a CPE's basic connection ID as done when performing an individual poll, the base station associates the allocated bandwidth to a multicast or broadcast connection ID. This multicast/broadcast bandwidth allocation is shown in the multicast/broadcast uplink sub-frame map **400"** shown in FIG. **7**. It is instructive to compare the uplink subframe **400** (FIG. **4**) used by the base station when individual polling the CPEs with the uplink sub-frame map **400"** of FIG. **7**. FIG. **7** shows the uplink sub-frame map which is transmitted in the MAC control portion of the downlink.

As shown in FIG. **7**, the multicast/broadcast uplink subframe map **400"** used by the present invention includes registration contention slots **402"** that map the registration contention slots **402** of FIG. **4**. However, rather than associating allocated bandwidth with a selected CPE's basic connection ID, the allocated bandwidth is associated with a reserved registration ID. As shown in FIG. **7**, the uplink sub-frame map **400"** preferably includes a plurality of multicast group bandwidth request contention slots **404"**, **404'"**, etc. The uplink sub-frame map **400"** also includes broadcast bandwidth request contention slots **410**. Finally, similar to the uplink sub-frame of FIG. **4**, the uplink sub-frame map used by the present invention to initiate multicast or broadcast polls includes a plurality of CPE scheduled data blocks **406"**, **406'"**, etc., that are used to transport uplink traffic data.

In accordance with the present inventive method and apparatus, when a poll is directed to a multicast or broadcast connection ID, CPEs belonging to the polled group request bandwidth using the bandwidth request contention slots (either the multicast contention slots for the group specified or the broadcast bandwidth request contention slots **410**) allocated in the uplink sub-frame map **400"**. In order to reduce the likelihood of collisions only CPE's needing bandwidth are allowed to reply to multicast or broadcast polls. Zero-length bandwidth requests are not allowed in the bandwidth request contention slots. In one embodiment, CPEs transmit the bandwidth requests in the bandwidth request contention slots (e.g., contention slots **404**) using QAM-4 modulation. In this embodiment, the contention slots are sized to hold a 1-PS preamble and a bandwidth request message. Due to physical resolution characteristics, the message requires 1 PI (or 6 PS) using QAM-4 modulation. In this embodiment, multiple bandwidth request messages from the same CPE fit in a single bandwidth request contention slot without increasing the bandwidth utilization or the likelihood of collisions occurring. This allows the same CPE to make multiple bandwidth requests in the same slot.

If an error occurs when performing either a multicast or broadcast poll (such as the detection of an invalid connection ID) the base station transmits an explicit error message to the CPE. If the base station does not respond with either an error message or a bandwidth allocation within a predefined time period, the CPE will assume that a collision occurred. In this case the CPE uses a selected pre-defined contention resolution process. For example, in one preferred embodiment, the CPE uses the well-known "slotted ALOHA" contention resolution process to back off and try at another contention opportunity.

Contention Resolution Process

Contention is necessary when there is insufficient time to poll all of the CPEs individually within a suitable interval. The base station is able to define contention periods both for multicast groups and also for all CPEs generally (i.e., broadcast). After CPE scheduled data, control messages, and polling are allowed for, the base station allocates all unused time in the upstream part of the TDD frame to contention, either for bandwidth requests or for registration purposes. Typically the bandwidth request interval will be many PIs long (e.g., 1 PI=6 PS using QAM-4 modulation). The CPEs must transmit their requests at a random time (on burst boundaries) within this interval to reduce the likelihood of collisions occurring.

In accordance with the present invention, a CPE needing to transmit in a request interval preferably randomly selects a PI within the interval, and makes a request in the associated starting PS. This randomization minimizes the probability of collisions. A collision is presumed if there is no response from the base station to the request within a pre-defined time period. If the base station does not respond within the predefined time period the collision resolution process of the present invention is initiated.

One preferred embodiment of the present invention uses the following resolution process: Assuming that the initial backoff parameter is i and that the final backoff parameter is f,

1. On the first collision, the CPE waits a random interval between zero and 2.sup.i contention opportunities and then tries again.

2. If another collision occurs, then the interval is doubled and the CPE tries again, repeating until the interval 2.sup.f is reached.

If the CPE is still unsuccessful, an error is reported to the system controller and the contention process is aborted. Other contention resolution mechanisms can be used to practice the present invention. For example, the well-known Ternary tree mechanism could be used to resolve contentions.

FIG. **8** is a flowchart showing the multicast and broadcast polling method **800** of the present invention. As shown in FIG. **8**, the group polling method **800** proceeds from an initial step **802** to a decision STEP **804** whereat the method determines whether there is sufficient bandwidth available for multicast polls. If sufficient bandwidth is available for multicast polls, the method proceeds to a STEP **806** to poll the next multicast group in the MAC control portion **314** of the MAC frame control header **302**. However, if there is insufficient bandwidth available to perform a multicast poll, the method proceeds to a decision STEP **808** whereat the method determines whether there is sufficient available bandwidth for performing a broadcast poll. If so, the method proceeds to a STEP **810**. If not, the method proceeds to a decision STEP **812**.

As shown in FIG. **8**, at the STEP **810** a broadcast poll is initiated by placing the broadcast poll in the MAC control portion **314** of the MAC frame control header **302**. Similar to the individual polling technique, the multicast poll message is implicitly transmitted to the CPE by allocating bandwidth in the uplink sub-frame map **400"**. The allocated bandwidth is associated with a multicast or broadcast connection ID.

At the decision STEP **812** the method determines whether a broadcast or multicast poll was initiated. If so, the method proceeds to a STEP **814** whereat the method monitors the appropriate bandwidth request contention slots (e.g., as defined by the bandwidth contention slot descriptions **404"**, **404'"**, and the broadcast bandwidth request contention slot descriptions **410** of FIG. **7**). If no broadcast or multicast poll

was initiated, the method proceeds to control point **816** and subsequently terminates at a termination STEP **818**.

The method proceeds from the monitoring STEP **814** to a decision STEP **820** to determine whether valid (i.e., non-colliding) bandwidth requests were detected. If no valid bandwidth requests were detected at STEP **820**, the method proceeds to the control point **816** and terminates at termination STEP **818**. However, if the method detects valid bandwidth requests, the method proceeds from STEP **820** to STEP **822**. At STEP **822** the method uses a convenient bandwidth allocation algorithm to allocate bandwidth to the CPE that requested bandwidth. The preferred bandwidth allocation algorithm is described below in more detail with reference to FIGS. **12-13**. The bandwidth is allocated in the uplink sub-frame map **400"** as shown in FIG. **8**.

Poll-Me Bit

As described above with reference to FIGS. **3-8**, and in accordance with the present invention, a currently active CPE sets a "poll-me" bit or a "priority poll-me" in a MAC packet in order to indicate to the base station that it requires a change in bandwidth allocation. For example, in one embodiment of the present invention, a selected CPE requests a poll by setting a poll-me ("PM") bit in the MAC header. Similarly, in accordance with the present invention, a selected CPE sets a priority poll-me ("PPM") bit in the MAC header in order to indicate that a priority poll is desired.

In order to reduce the bandwidth requirements associated with individually polling every active CPE, the active CPEs are individually polled if and only if one of the poll-me bits is set by the CPE. When the base station detects a request for polling (when the CPE sets its poll-me bit), the individual polling technique shown in FIG. **9** is activated in order to satisfy the request. The procedure by which a CPE stimulates a base station to poll the CPE is shown in FIG. **9**. In an alternative embodiment, multiple packets having "poll-me" bits set indicate that the CPE needs to make bandwidth allocation requests for multiple connections.

FIG. **9** is a flow chart that shows how the poll-me bit is used to stimulate polling in accordance with the present invention. As shown in FIG. **9**, the method first determines at a decision STEP **902** whether the piggybacking technique described in more detail below has been exhausted. If not, the method proceeds to STEP **904** and attempts to perform "piggybacking" first. The method then proceeds to a STEP **906** whereat the connection is set equal to a first connection. In this manner, the poll-me bits are scanned for each connection within the CPE. The method shown in FIG. **9** then proceeds to a decision STEP **908** to determine whether any bandwidth needs exist. If not, the method proceeds to a STEP **916** and scans for the next connection. If a bandwidth need exists, the method proceeds to a decision STEP **910**. At STEP **910** the method determines whether any more packets are available for accommodating the poll-me bit. If not, the method terminates at the STEP **910**. However, if packets are available, the method proceeds to a STEP **912** and sets a poll-me bit in an available packet.

FIG. **10** shows the message sequence that is used by the present invention in requesting polls using the "poll-me" bit described above. As shown in FIG. **10** at data connection **930**, the CPE initiates a polling sequence by setting its associated poll-me bit in the MAC header. The base station MAC responds via data message **932** by individually polling the selected CPE. This response is made by allocating bandwidth to the selected CPE in the uplink sub-frame map as shown in FIG. **10**. The selected CPE subsequently responds with a bandwidth request as shown in communication path **934**. In response to the CPE's bandwidth request, the base station grants bandwidth and allocates bandwidth to the CPE in the uplink sub-frame map as shown in communication path **936**. The selected CPE then transmits its data to the base station via an associated connection link.

"Piggybacking" Technique

As described above with reference to the present inventive method and apparatus, in order to further reduce overhead bandwidth necessary for the bandwidth allocation process, currently active CPEs may "piggyback" a bandwidth request (or any other control message) on their current transmissions. The CPEs accomplish this piggybacking of bandwidth by using unused bandwidth in TC/PHY packets of existing bandwidth allocations. The procedure for using excess bandwidth in this manner is shown in FIG. **11**.

As shown in FIG. **11**, the method initiates the piggybacking process at STEP **950**. The method proceeds to a decision STEP **952** to determine whether the CPE requires additional bandwidth. If so, the method proceeds to a decision STEP **954**, if not, the method proceeds to a termination STEP **964** whereat the method terminates. At the decision STEP **954** the method determines whether any unused bytes exist in the current allocation. If so, the method proceeds to insert bandwidth requests into the unused bytes at STEP **956**. If not, the method proceeds to a decision STEP **958**. At the decision STEP **958**, the method determines whether any packets at all are allocated to the CPE. If there are no packets found at the decision STEP **958**, the method proceeds to STEP **960**. However, if packets are allocated, the method proceeds to a STEP **962** whereat the CPE sets its poll-me bit. The method then proceeds to the STEP **960** whereat the CPE awaits polling by the associated base station. The method then terminates at the STEP **964**.

Bandwidth Allocation

As described above, the base station MAC is responsible for allocating the available bandwidth of a physical channel on the uplink and the downlink. Within the uplink and downlink sub-frames, the base station MAC scheduler allocates the available bandwidth between the various services depending upon the priorities and rules imposed by their quality of service (QoS). Additionally, the higher control sub-layers of the base station MAC allocate across more than one physical channel.

Down link Bandwidth Allocation—One Embodiment

The downlink bandwidth is allocated as shown in FIG. **12**. The base station MAC maintains a set of queues for each physical channel that it serves. Within each physical channel queue set, the base station maintains a queue for each QoS. The queues hold data that is ready to be transmitted to the CPEs present on the physical channel. The higher layers of the base station protocol stack are responsible for the order in which data is place in the individual queues. The base station higher control layers are free to implement any convenient fairness or traffic shaping algorithms regarding the sharing of access between connections at the same QoS, without impacting the base station lower MAC control layers. Once data is present in the queues it is the responsibility of the base station lower levels of control (e.g., the BS LL-MAA of FIGS. **5** and **10**) to allocate bandwidth based on the QoS.

In one embodiment of the present invention, in determining the amount of bandwidth to allocate at a particular QoS for a particular CPE, the base station takes into account the QoS, modulation, and the fairness criteria used to keep an individual CPE from using up all available bandwidth. Bandwidth is preferably allocated in QoS order. If there is a queue that cannot be transmitted entirely within a particular TDD frame, a QoS specific fairness algorithm, such as fair-weighted queuing, is used within that queue. Each connection

is given a portion of the remaining available bandwidth based upon its relative weight. The derivation of weights is QoS-dependant. For example, ATM traffic may be weighted based upon contractual bandwidth limits or guarantees, while IP connections may all receive identical weights. Once the bandwidth is allocated the data is transmitted in a manner whereby the data is sorted by modulation type.

Uplink Bandwidth Allocation—One Embodiment

The uplink bandwidth allocation method is very similar to the downlink bandwidth allocation method described above with reference to FIG. **12**. However, rather than being maintained by the base station, the data queues are distributed across and maintained by each individual CPE. Rather than check the queue status directly, the base station preferably receives requests for bandwidth from the CPEs using the techniques described above with reference to FIGS. **3-11**. Using these bandwidth requests, the base station reconstructs a logical picture of the state of the CPE data queues. Based on this logical view of the set of queues, the base station allocates uplink bandwidth in the same way as it allocates downlink bandwidth. This uplink bandwidth allocation technique is shown in FIG. **13**.

As described above, the bandwidth allocated to any selected CPE is transmitted to the selected CPE in the form of bandwidth being allocated in the uplink sub-frame map. Starting at a point in the TDD, the uplink sub-frame map allocates a certain amount of bandwidth to the selected CPE. The selected CPE then allocates this bandwidth across its connections. This allows the CPE to use the bandwidth in a different manner than requested if it receives higher priority data while awaiting the bandwidth allocation. As described above, the bandwidth allocations are in a constant state of change owing to the dynamic nature of bandwidth requirements. Consequently, a selected CPE may receive unsolicited modifications to the bandwidth granted on a frame-by-frame basis. If the selected CPE is allocated less bandwidth for a frame than is necessary to transmit all waiting data, the CPE must use the QoSs and fairness algorithms to service its queues. The CPE may “steal” bandwidth from lower QoS connections to piggyback request for more bandwidth using the piggybacking technique described above. TDM connections not already at maximum bandwidth are allocated enough extra bandwidth in the uplink to piggyback a request for additional bandwidth.

QoS Specific Fairness Algorithms

Data for transmission on the uplink and the, downlink is preferably queued by quality of service (QoS) designations. The data is transmitted in order of a QoS queue priority as described above. As the queued data is transmitted, there may be a QoS queue for which there is insufficient bandwidth to transmit all queued data during the current TDD frame. When this situation occurs, a QoS specific fairness algorithm is initiated to ensure fair handling of the data queued at that QoS. There are 3 basic fairness algorithms that can be implemented: (1) Continuous Grant; (2) Fair-weighted queuing; and (3) Round Robin.

The MAC preferably does not police connections for bandwidth usage. Policing should be performed by higher control layers. The MAC assumes that all pending data has met contractual restrictions and can be transmitted. Continuous Grant queues have the simplest fairness algorithm. All data in these queues must be sent every TDD frame. Insufficient bandwidth indicates an error in provisioning.

Fair Weighted Queuing

Fair weighted queuing requires that all connections at a given QoS have a weight, assigned to them to determine the percentage of the available bandwidth they are eligible to receive. This weight value is preferably derived from one of three data rate parameters, depending upon the contractual parameters of the provisioned connection. These three parameters are: (1) Data Pending; (2) Guaranteed Rate; and (3) Average Rate.

Real-time VBR connections are established as DAMA connections with fair-weighted queuing based upon data pending. For a QoS queue of this type in a TDD frame having insufficient bandwidth to transmit all of the data in the queue, a weight for each connection in the queue is determined. In one embodiment, this weight is the amount of data pending for the connection expressed as a percentage of the total data pending in the queue. Because the amount of data pending is dynamic, the weights for these types of queues must be determined every TDD frame where there is insufficient bandwidth to send all data in the affected queue.

For DAMA connections contracted at a guaranteed rate the weights are calculated based on the guaranteed rate. In this case, the weight preferably is expressed as a percentage of the total guaranteed rate of all connections with data pending in the queue. Because the guaranteed rate is provisioned the weights need not be determined each TDD frame where they are used. Rather, the weights for a queue are only determined when there is a provisioning change (i.e., a new connection, a change in connection parameters, or a connection termination) for one of the connections in the queue.

For DAMA connections contracted at an average rate the weights are preferably calculated based on the average rate. The weight is the average rate expressed as a percentage of the total average rate of all connections with data pending in the queue. Because the average rate is provisioned the weights need not be determined each TDD frame where they are used. Rather, the weights for a queue are only recalculated when there is a provisioning change for one of the connections in the queue.

In all of the cases described above, the granularity of the bandwidth allocations may be too coarse to provide a perfect percentage-based weighted allocation across the connections in the queue. This may result in some queues not receiving any bandwidth in a particular TDD frame. To ensure that the occurrence of this condition is fairly distributed across the connections in the queue, the connection that did not receive bandwidth is given priority the next time the insufficient bandwidth condition exists for the queue. For queues with weights based upon guaranteed or average rates some connections may not have sufficient data pending to use all of the bandwidth that they are entitled to based upon their calculated weight. In these cases, the connection’s unused bandwidth is fairly distributed across the connections having excess data pending.

Some QoSs require that data be aged. For queues at these QoSs there is an associated queue of one step higher priority. If data is not transmitted by the provisioned aging parameter, the data is moved to the higher QoS queue and given priority over newer data in the original queue regardless of the relative weights of the connections.

Round Robin

The Round Robin fairness algorithm is used for best effort connections where all connections have equal weight. When insufficient bandwidth exists to transmit all data in the queue in a particular TDD frame connections are allocated bandwidth in a round-robin fashion with each connection receiving a block of bandwidth up to a queue-specific maximum. Connections that did not receive bandwidth are given priority the next time the insufficient bandwidth condition exists.

Bandwidth Allocation Algorithm

For each TDD frame, the base station allocates the downlink portion of the TDD frame and it performs an estimate of the uplink traffic to allocate uplink bandwidth to the CPEs. The CPEs individually allocate their allotted bandwidth across their pending data connections.

Base Station Downlink

As shown in FIG. **2**, in one preferred embodiment of the present invention, based on the ATDD split (i.e., the percentage of bandwidth allocated to the uplink and downlink) the base station has some number of the 800 PS in the TDD frame available for downlink transmissions. The downlink bandwidth allocation algorithm preferably proceeds as follows.

First, the base station allocates PSs to the PI for PHY Control and enough PSs for at least 1 PI for the MAC Control. The base station preferably performs uplink bandwidth allocation before downlink bandwidth allocation in order to determine the number of PIs to allocate for the MAC Control. In one preferred embodiment, the PHY Control and MAC Control are always sent using QAM-4 modulation.

For connections with downlink continuous grant data pending, the base station determines the number of PIs required to transmit the data. This number is then converted to PSs as a function of the modulation used for the CPE associated with each connection. For each remaining QoS or until available bandwidth is entirely allocated, the base station determines if there is enough bandwidth to satisfy the entire need of the QoS queue. If so, the base station allocates the required bandwidth. Otherwise, if there is not enough bandwidth to satisfy the queue, the base station implements the queue-specific fairness algorithm described above.

Base Station Uplink

In one preferred embodiment, based upon the ATDD split described above with reference to FIG. **2**, the base station has a pre-determined number of PSs in the TDD frame available for uplink transmissions. The base station must maintain an estimate of the data and control messages pending at each QoS for the CPEs that it serves. The base station estimates the data traffic based upon the bandwidth requests received from the CPEs and based upon an observation of actual data traffic. The base station estimates the uplink control message traffic based upon the protocols currently engaged (i.e., connection establishment, "poll-me" bit usage, etc.) and based upon the base station's polling policy (i.e., individual, multicast, and broadcast). The uplink bandwidth allocation algorithm proceeds as follows.

For connections with uplink continuous grant data pending, the base station preferably determines the number of PIs required to transmit the data. This number is then converted to a number of PSs as determined by the modulation used for the CPE associated with each connection. Continuous grant connections having a current bandwidth that is less than the maximum bandwidth are always allocated uplink bandwidth that is the smaller of: 1) their maximum bandwidth or 2) their current bandwidth plus the bandwidth necessary to send a CG bandwidth change message.

For each remaining QoS, or until available bandwidth is entirely allocated, the base station determines if there is bandwidth sufficient to satisfy the entire need of the QoS queue and it then allocates the required bandwidth. Otherwise, if there is not bandwidth sufficient to satisfy the queue, the base station implements the queue-specific fairness algorithm described above.

CPE Uplink

As described above, for each TDD frame, the CPEs are allocated a portion of the uplink sub-frame in which to transmit their respective data. Because the bandwidth requirements of the CPE may have changed since the base station received the bandwidth request information that it used to allocate the uplink bandwidth, the CPEs themselves are responsible for allocating their allotted bandwidth based upon their current bandwidth requirements. That is, the CPEs are not constrained to distribute allocated bandwidth to their data connections in the same manner that the CPE used in requesting the bandwidth from the base station. The CPE's uplink bandwidth allocation algorithm preferably proceeds as follows.

For connections having uplink continuous grant data pending, the CPE determines the number of PIs that are required to transmit the data. This number is then converted to a PS number based upon the modulation scheme used by the CPE. For each remaining QoS, or until available bandwidth is entirely allocated, the CPE determines if there is bandwidth sufficient to satisfy the entire need of the QoS queue. If so, the CPE allocates the required bandwidth. Otherwise, if there is not bandwidth sufficient to satisfy the queue, the CPE implements the queue-specific fairness algorithm described above.

SUMMARY

In summary, the bandwidth allocation method and apparatus of the present invention includes a powerful, highly efficient means for allocating bandwidth in a broadband wireless communication system. The present bandwidth allocation method and apparatus uses a combination of individual and group polling techniques, contention-based polling, piggybacking, and CPE-initiated techniques to efficiently allocate bandwidth in a communication system. Advantageously, only those currently active CPEs (CPEs that currently have bandwidth allocations associated thereto) are permitted to request more bandwidth using either the piggybacking or poll-me bit methods. In addition, the present invention saves bandwidth by implicitly informing the CPE of additional bandwidth allocation. The base station implicitly informs the CPE of additional bandwidth allocation by allocating additional bandwidth to the CPE in the uplink sub-frame map. Similarly, the base stations implicitly poll the CPEs by allocating bandwidth in the uplink to enable the CPEs to respond to the poll with a bandwidth request.

In honoring the bandwidth requests, the base station builds and maintains a logical queue of the data to be transmitted. The queues are developed by the base stations based upon the QoS. In addition, the base station allocates bandwidth based on a combination of QoS and a QoS unique fairness algorithm. The CPE itself, rather than the base station, distributes the allocated bandwidth to its services in any manner the CPE determines to be appropriate. Thus, the CPE can use its allocated bandwidth in a manner that differs from the originally intended (and requested) purpose.

A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, the present inventive method and apparatus can be used in any type of communication, its use is not limited to a wireless communication system. One such example is use of the invention in a satellite communication system. In such a communication system, satellites replace the base stations described above. In addition, the CPEs are not longer at fixed distances from the satellites. Therefore, it will be more difficult to schedule DAMA services for the CPEs. Alternatively, the present invention can be used in a wired communication system. The only difference between the wired system and the wireless system described above is that the channel characteristics

vary between the two. However, the bandwidth allocations do not change as between the two types of systems. Accordingly, it is to be understood that the invention is not to be limited by the specific illustrated embodiment, but only by the scope of the appended claims.

The invention claimed is:

1. A method for requesting bandwidth on demand in a wireless communication system, wherein the wireless communication system includes a wireless subscriber radio unit, the method comprising:

registering the wireless communication radio unit with a base station in the wireless communication system and establishing communication between the wireless subscriber radio unit and the base station;

transmitting from the wireless subscriber radio unit which is registered with the base station, an explicit message to the base station requesting to be provided an allocation of uplink (UL) bandwidth in which to transmit a bandwidth request;

receiving at the wireless subscriber radio unit the allocation of UL bandwidth in which to transmit a bandwidth request;

transmitting the bandwidth request within the allocation of UL bandwidth, the bandwidth request specifying a requested UL bandwidth allocation; and

receiving an UL bandwidth grant for the wireless subscriber radio unit in response to the bandwidth request;

wherein the wireless subscriber radio unit maintains a plurality of queues, each queue for data pertaining to one or more UL connections with similar QoS and wherein the wireless subscriber radio unit allocates the UL bandwidth grant to the one or more UL connections based on QoS priority.

2. A method as claimed in claim **1**, wherein the bandwidth request is transmitted by the wireless subscriber radio unit within UL bandwidth specifically allocated to the subscriber unit for the bandwidth request.

3. A method as claimed in claim **1**, wherein the wireless subscriber radio unit contends with other wireless subscriber radio units for transmitting the bandwidth request within UL bandwidth specifically allocated to a group of wireless subscriber radio units for bandwidth requests.

4. A method as claimed in claim **1**, wherein the bandwidth request pertains to a connection.

5. A method as claimed in claim **1**, further comprising the wireless subscriber radio unit allocating the UL bandwidth grant to at least one of the plurality of UL connections established at the wireless subscriber radio unit.

6. A method of allocating uplink (UL) bandwidth on demand in a wireless subscriber radio unit, comprising:

registering the wireless subscriber radio unit with a base station in a wireless communication system and establishing communication between the wireless subscriber radio unit and the base station;

determining a required amount of UL bandwidth for a UL queue at the wireless subscriber radio unit, wherein the UL queue comprises traffic with similar quality of service (QoS) received on a plurality of connections;

transmitting from the wireless subscriber radio unit an explicit message to the base station requesting to be provided an allocation of UL bandwidth in which to transmit a bandwidth request;

receiving at the wireless subscriber radio unit the allocation of UL bandwidth in which to transmit a bandwidth request;

sending the bandwidth request indicative of the required amount of UL bandwidth for the UL queue to the base station;

receiving from the base station an UL bandwidth grant for the wireless subscriber radio unit; and

allocating, at the wireless subscriber radio unit, the UL bandwidth grant to one or more of a plurality of connections at the wireless subscriber radio unit;

wherein the wireless subscriber radio unit allocates the UL bandwidth grant to the one or more of the plurality of connections based on QoS connection priority.

7. A method as claimed in claim **6**, wherein the bandwidth request is transmitted by the wireless subscriber radio unit within UL bandwidth specifically allocated to the wireless subscriber radio unit for the bandwidth request.

8. A method as claimed in claim **6**, wherein the wireless subscriber radio unit contends with other wireless subscriber radio units for transmitting the bandwidth request within UL bandwidth specifically allocated to a group of wireless subscriber radio units for the bandwidth request.

9. A method as claimed in claim **6**, wherein the bandwidth request is transmitted by the wireless subscriber radio unit within UL bandwidth already allocated to the wireless subscriber radio unit for UL traffic.

10. A method as claimed in claim **6**, wherein the bandwidth request is transmitted by the wireless subscriber radio unit in unused UL bandwidth allocated to the wireless subscriber radio unit.

* * * * *

CERTIFICATE OF SERVICE

I, Daniela Lattes, hereby certify that on March 13, 2015, I caused one copy of the documents listed below:

NON-CONFIDENTIAL BRIEF FOR PLAINTIFFS-APPELLANTS WI-LAN USA INC. AND WI-LAN, INC.

to be filed by CM/ECF with:

Clerk of Court
United States Court of Appeals for the Federal Circuit
717 Madison Place, N.W.
Washington, D.C. 20439
Tel: (202) 275-8000
Fax: (202) 275-9678

and to be served via CM/ECF and electronic mail pursuant to Fed. Cir. R. ECF-9(b) on the following:

Mark C. Scarsi
mscarsi@milbank.com
Christopher J. Gaspar
cgaspar@milbank.com
Miguel Ruiz
mruiz@milbank.com
Jennifer L. Miremadi
jmiremadi@milbank.com
Ashlee N. Lin
ashlee.lin@milbank.com
Michael Sheen
msheen@milbank.com
Milbank, Tweed, Hadley & McCloy, LLP
601 South Figueroa Street
Los Angeles, CA 90017
Phone : (213) 892-4459
Fax: (213) 892-4759

I declare that I am employed by McKool Smith P.C. at whose direction the service was made. Executed on March 13, 2015, at New York.

/*s/ Daniela Lattes*
Daniela Lattes
MCKOOL SMITH P.C.
One Bryant Park, 47th floor
New York, New York 10036
212.402.9499 (t)
212.402.9444 (f)
dlattes@mckoolsmith.com

CERTIFICATE OF COMPLIANCE

I certify that the foregoing NON CONFIDENTIAL OPENING BRIEF FOR PLAINTIFF-APPELLANT WI-LAN USA, INC. AND WI-LAN INC.:

1. Complies with the type-volume limitation of Fed. R. App. P. 32(a)(7)(B). This brief contains 8,209 words, excluding the parts of the brief exempted by Fed. R. App. P. 32(a)(7)(B)(iii) and Fed. Cir. R. 32(b). Microsoft Word 2003 was used to calculate the word count; and

2. Complies with the typeface requirements of Fed. R. App. P. 32(a)(5) and the type style requirements of Fed. R. App. P. 32(a)(6). This brief has been prepared in proportionally-spaced typeface using Microsoft Word 2003 in 14-point Times New Roman type style.

Dated: March 13, 2015

Respectfully submitted,

/s/ Robert A. Cote

Robert A. Cote
MCKOOL SMITH P.C.
1 Bryant Park
New York, NY 10036
(212) 402-9402

Attorney for Plaintiff-Appellant, Wi-LAN USA, Inc., and Wi-LAN Inc.